AMD RYZEN 7 9700X REVIEW PG. 74

PC OVERCLOCKING GUIDE PG. 32

BUILD A GAMING VPN PG. 68

MAXIMUM PC

MINIMUM BS • OCTOBER 2024 • www.maximumpc.com

WHY YOU MUST GO AMD

Build this super-powerful & reliable PC

- Outstanding stability and bang for buck
- 4K, 1080p, and AI performance tested
- **PLUS!** Budget 4TB SSDs are finally here

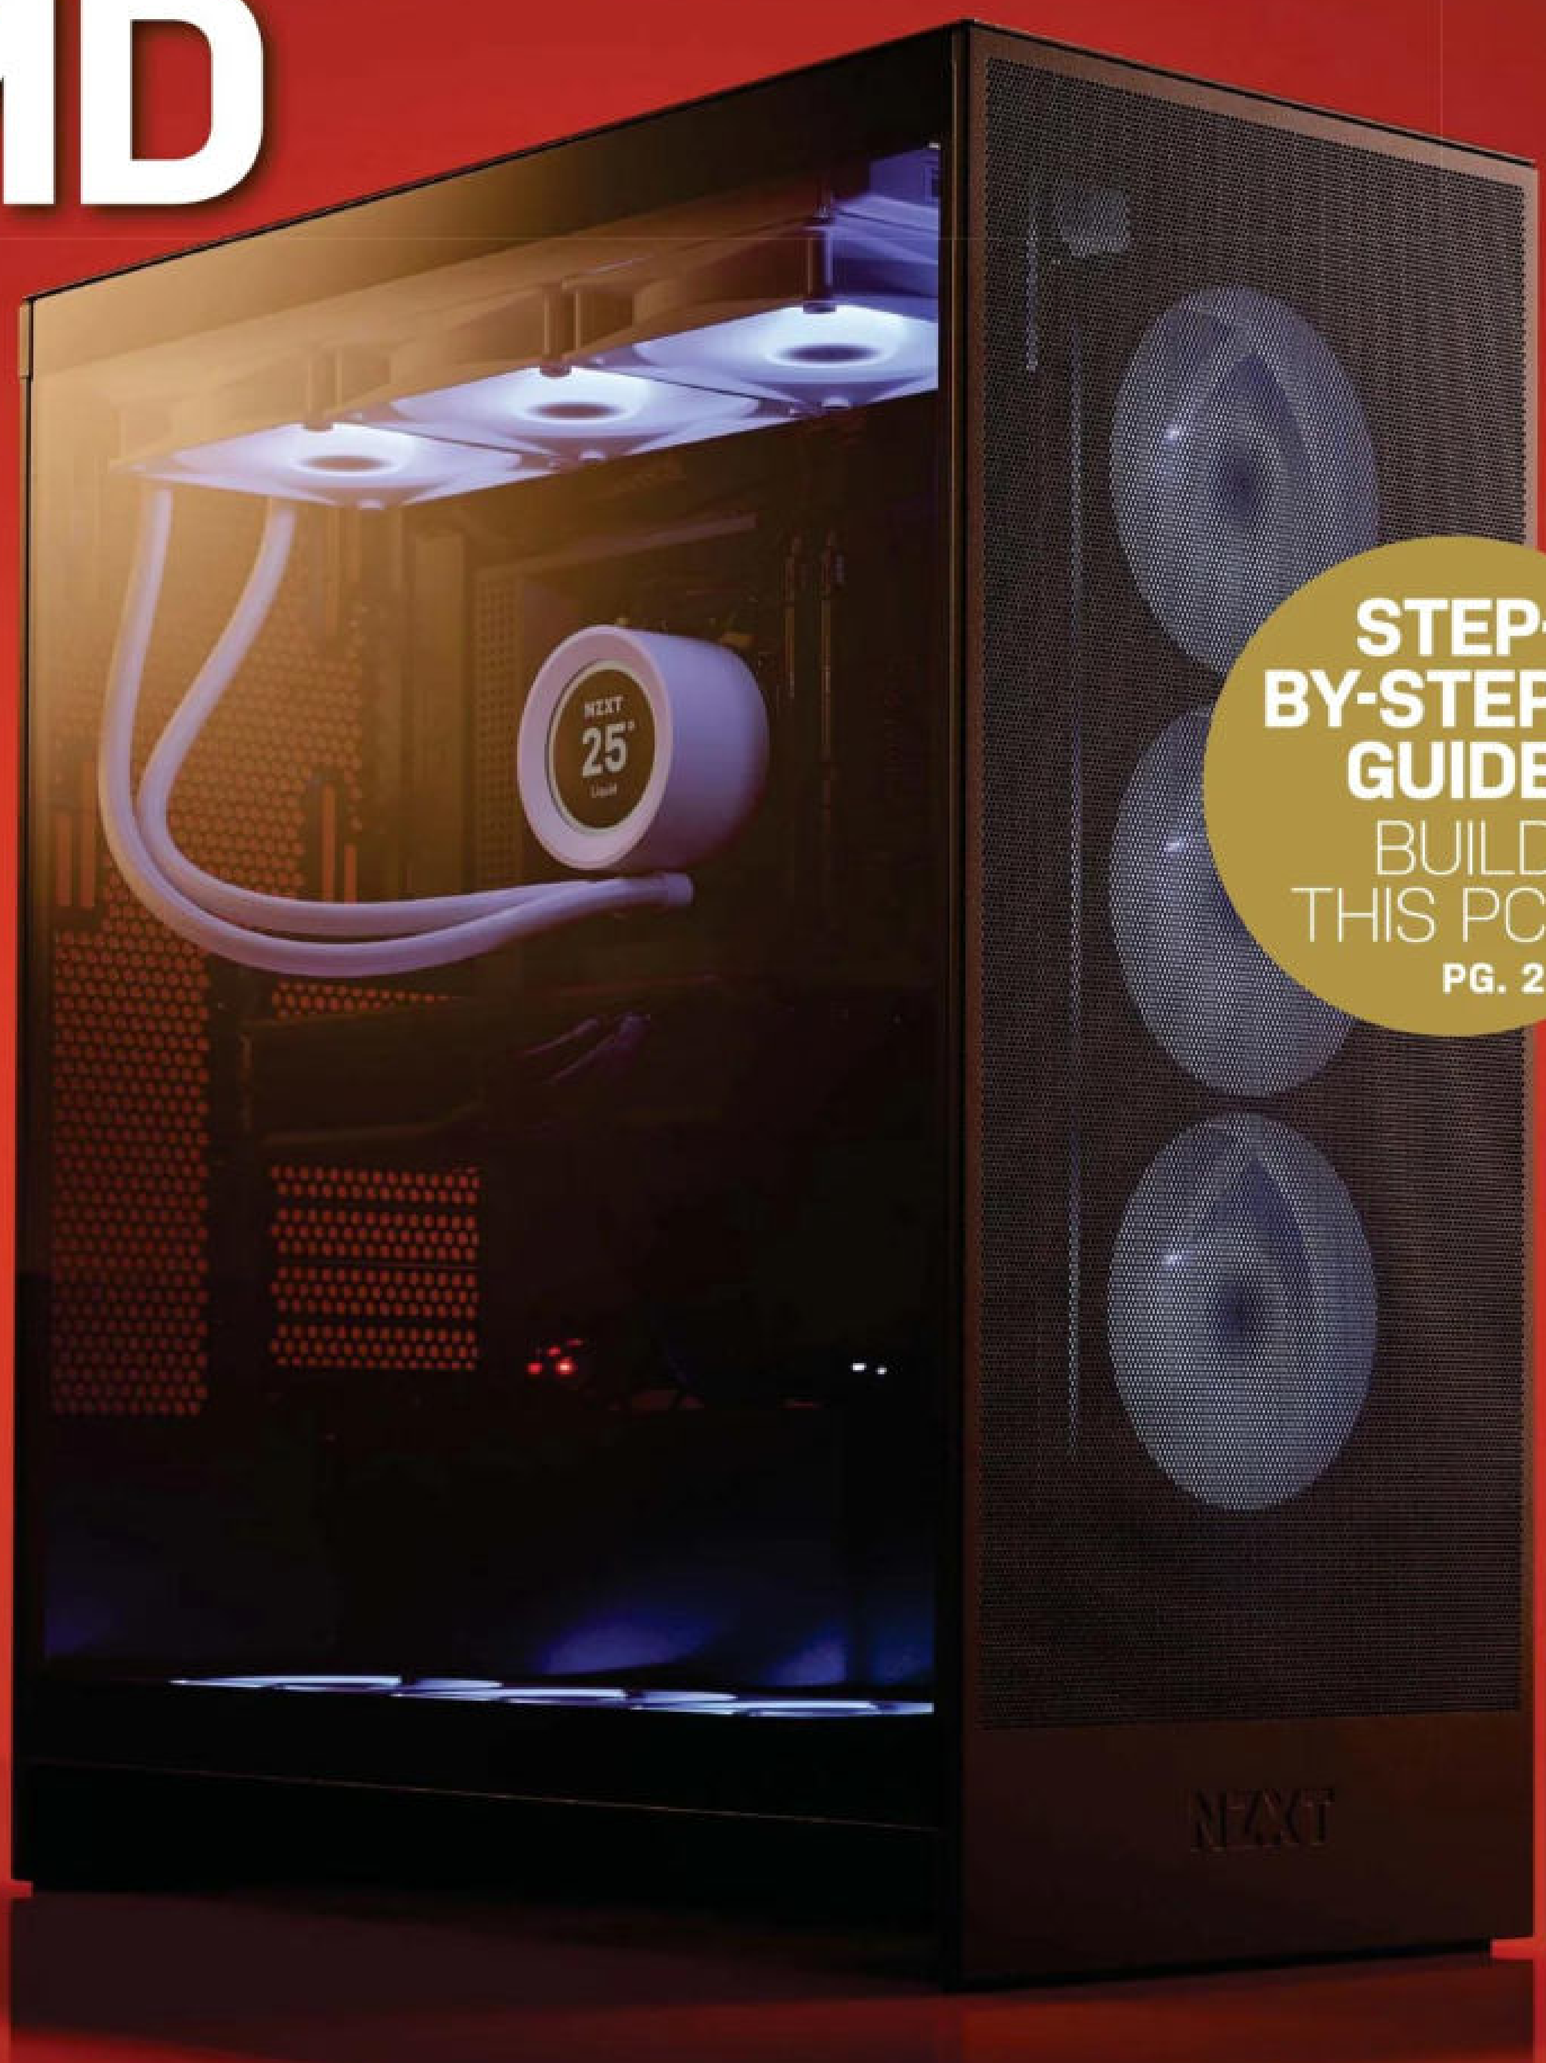

STEP-BY-STEP GUIDE
BUILD THIS PC!
PG. 21

NEXT-GEN CPUs EXPLAINED
PG. 50

HOW TO REPAIR YOUR GPU
PG. 70

FUTURE
VOL 29. NO 10


PC GAMER
EDITOR'S CHOICE
★★★★★
KICK ASS!
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 Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

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

OCTOBER 2024

SCAN TO GET THE TOM'S HARDWARE WEEKLY NEWSLETTER

QUICKSTART

R&D

LETTERS

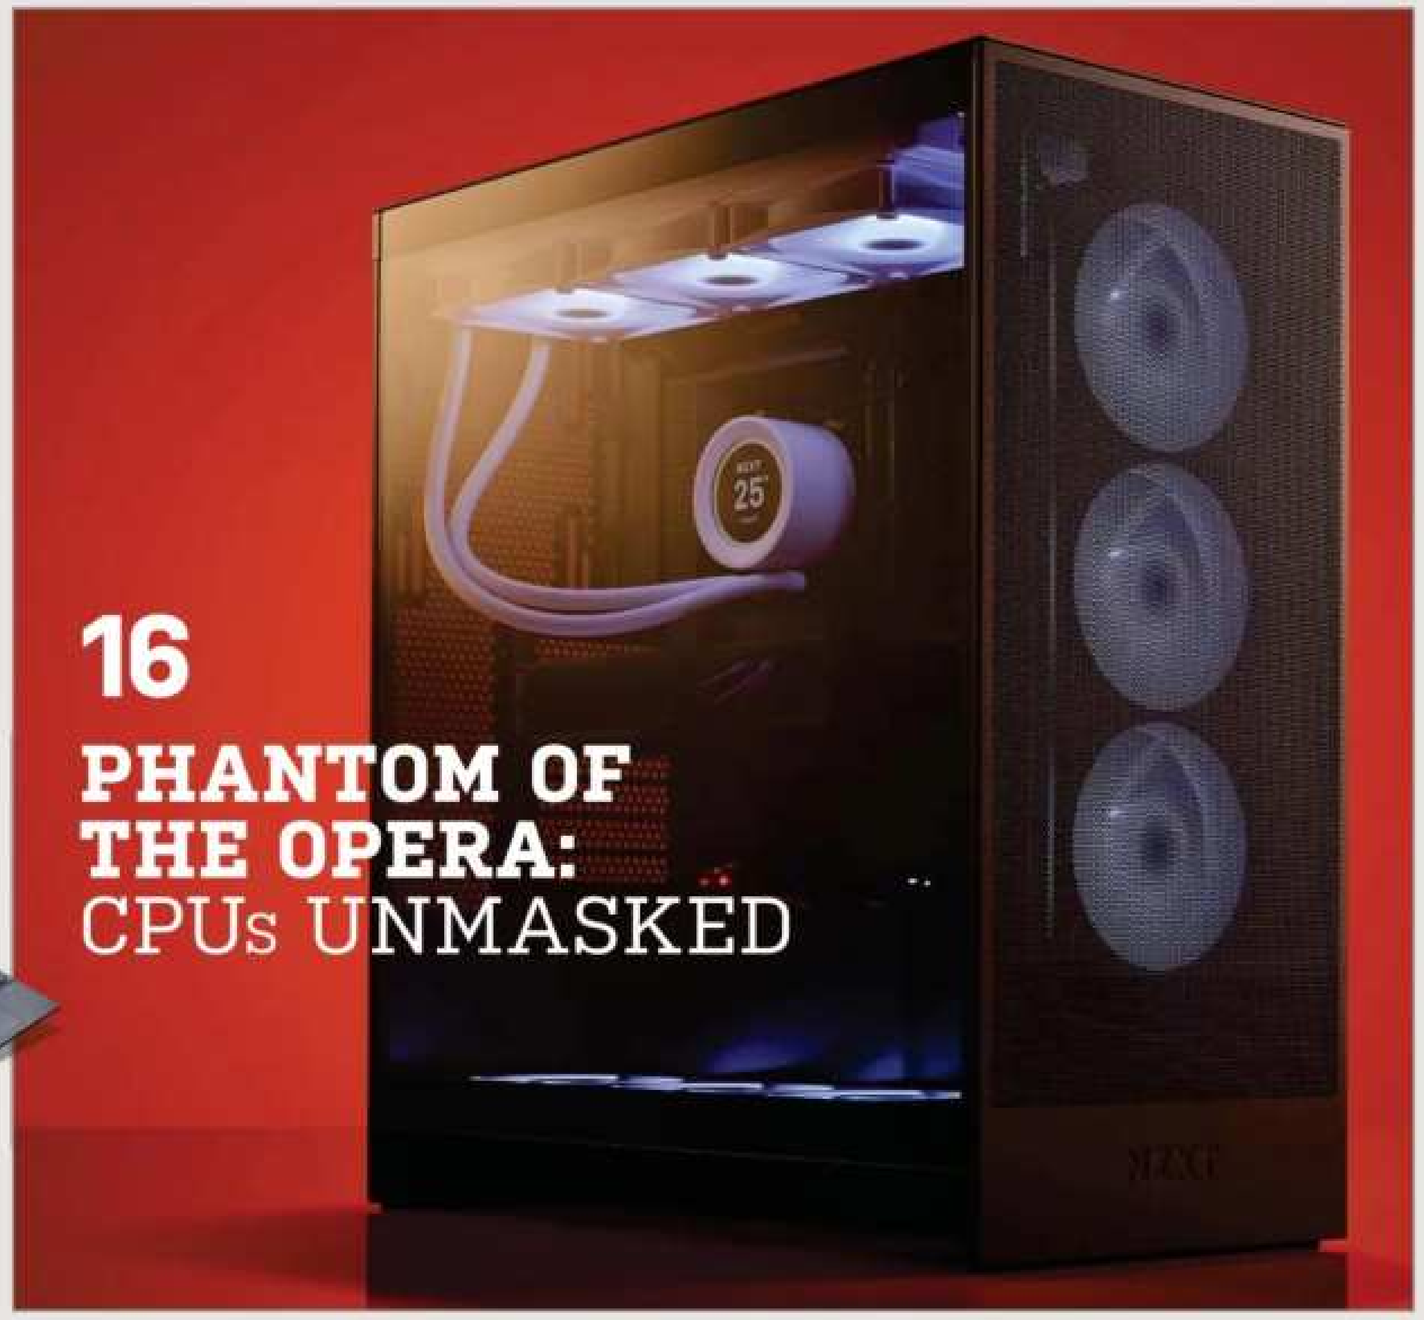

16
PHANTOM OF THE OPERA: CPUs UNMASKED

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

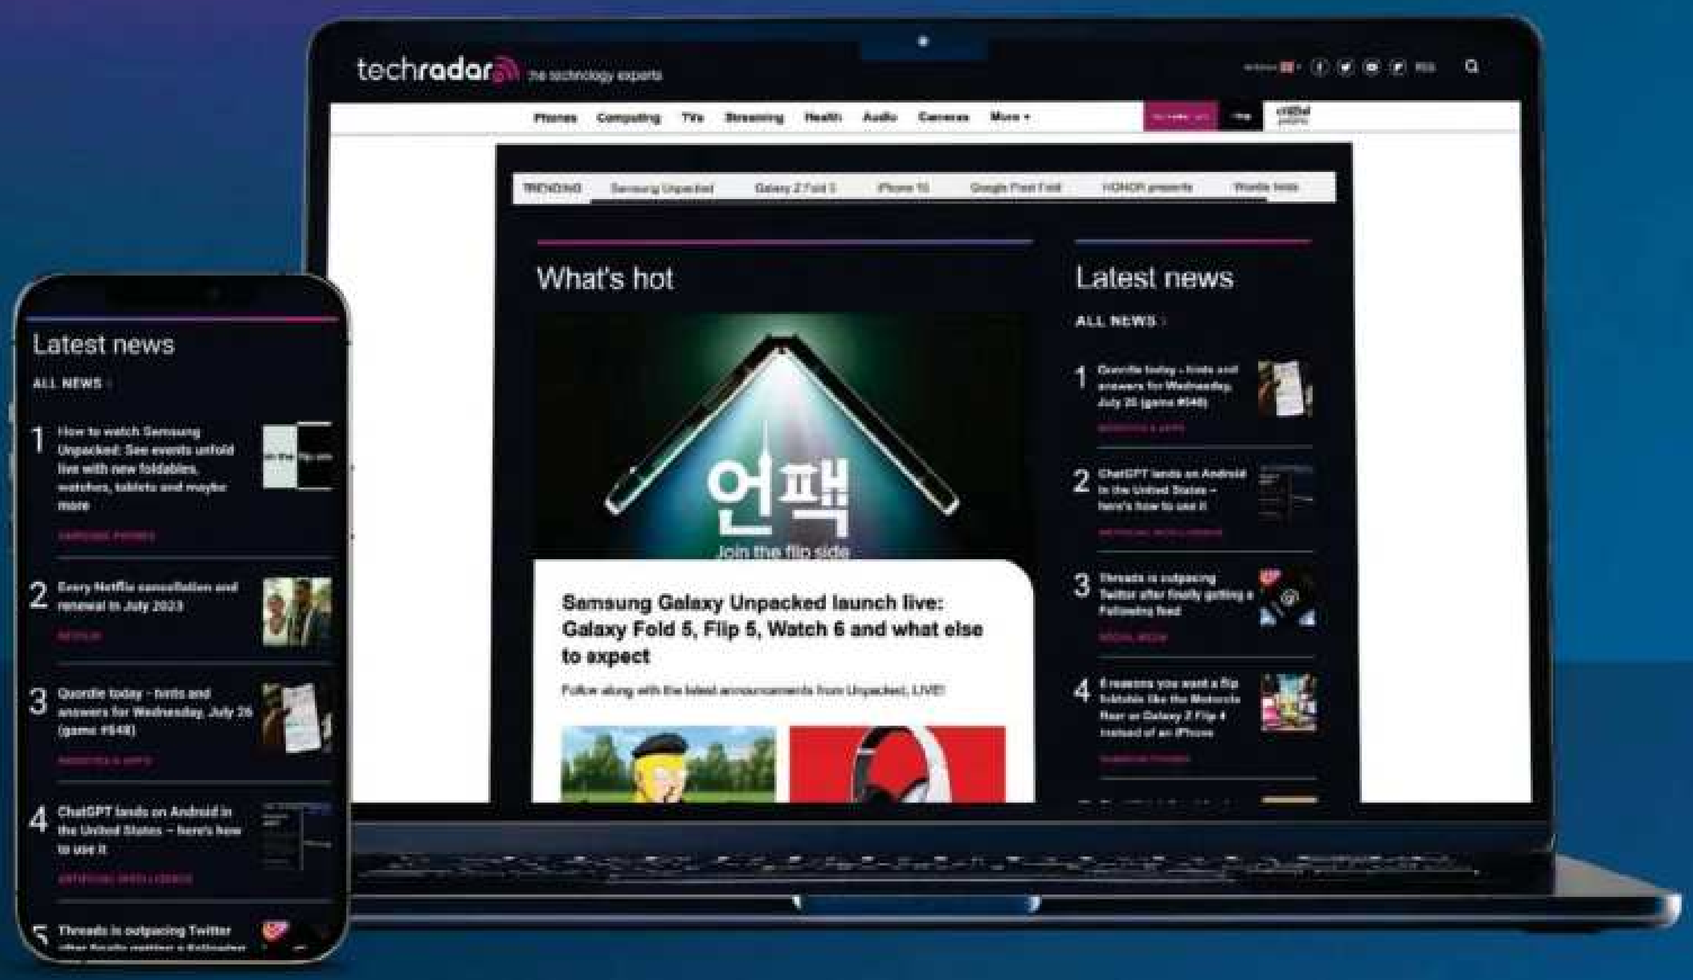

- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

techradar
the technology experts

www.techradar.com

MAXIMUMPC

EDITORIAL
Editor-in-Chief: Guy Cocker
Contributing Writers: Kerry Brunskill, Nate Drake, Nick Evanson, Ian Evenden, Emmanuel Fylladitakis, Jeremy Laird, Chris Lloyd, Nick Peers, Nik Rawlinson, Jacob Ridley, Zak Storey, Jarred Walton
Production Editor: Steve Wright
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Photography: Neil Godwin, Olly Curtis, Phil Barker
Cover Photo Credits: Getty Images, Scuf.com, 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 October 8, 2024

EXTRA DIGITAL FEATURES
AUDIO FILE | PHOTO GALLERY | VIDEO FILE

Guy Cocker

RELIABILITY IS KING

IT'S BEEN A STRANGE few months in the world of PC building. As you'll have read in these pages, Intel had been quiet about a crucial flaw that was discovered in its 13th and 14th-gen CPUs, a flaw that was so serious, some game developers said it was only a matter of time before all those chips became useless. Intel was fairly quiet on the subject, which resulted in us taking the affected models out of our Blueprint systems at the back of this magazine until we found out more.

Thankfully, Intel has acknowledged the issue, rolling out system updates to at least provide a temporary fix. They've also extended the warranty on these chips by two years, allowing users to get a replacement if they experience structural degradation. It's a bit of peace of mind for me, as I use an i5-14600K in my main system, but it still feels annoyingly like there's a ticking time bomb at the center of the desktop PC I use every day for work and gaming. I've never updated my motherboard BIOS as many times as I have in the last few months, as Intel rushes microcode fixes to manufacturers, who then roll them out as firmware updates.

Then, just as we were going to press, news hit of another vulnerability in millions of AMD chips that could compromise any systems using one, especially as it would be undetectable to virus and malware scanners. Thankfully, it's difficult to exploit, and AMD has again released firmware updates that fix the issue via a motherboard BIOS update. Sadly, AMD admits that some older chips won't be updated, as they're outside the company's support window. You can read about both issues in this month's news section, starting on page 8.

Given Intel's woes, I admit that I'd be more comfortable buying an AMD CPU right now, which is what's led to this issue's cover build. Our master builder, Zak, has attempted to create the most stable machine possible using an AMD 7900X 7th generation processor. The bonus? Given the next-generation 9000-series processors are just arriving (the first is our lead review on page 74), the 7900X build has the added benefit of being great value as well.

Zak's also taken this opportunity to update some of our benchmarks, including new AI-based ones for the first time. It all begins on page 16. Be sure to check out our video feature, Story of the Build, a link to which you'll also find on page 16. For those who like to be more on the cutting edge, worry not, as Zak should be back next issue with our first AMD 9000-series build.

Elsewhere, we take a look at overclocking on page 32—how to do it, what sort of results you can achieve, and whether it's worth doing in 2024. Speaking of CPUs, with the launch of a new generation from AMD and the next from Intel landing imminently, our expert Jarred Walton takes a deep dive into the next generation of processors, and what makes them worth the upgrade (or not). That's on page 50.

Finally, we have a whole host of laptop reviews, from a budget offering with fantastic performance, and a creator's laptop that's anything but boring, to the resurrection of the 18-inch gaming laptop with absolutely maxed-out specs—the sort of machine we used to call a 'desktop replacement'.

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

Intel's Raptor Lake liability

There's a fix for crashing 13th and 14th gen desktop CPUs

MONTHS AGO, some gamers in South Korea reported problems with their unlocked high-end Intel Core i9 processors, typically Core i9-13900K and 14900Ks. While playing *Tekken 8*, the game would shut down. More reports emerged involving other games, particularly titles using the Unreal engine. Chips were returned for replacement. Underclocking the chip slightly appeared to fix things. Motherboards were suspected of using overly aggressive voltage settings, while some claimed the chips weren't capable of sustaining the frequencies Intel stated. In April, Intel issued a statement that it was "conducting analysis of the reported issues".

In May, Intel issued another statement blaming the motherboard manufacturers. A BIOS update calmed down the default settings, which Intel claimed were not in line recommended defaults. Things improved, but the issue never went away. Then, in July, Intel released the result of its investigations on the returned chips. The problem was traced to "a microcode algorithm resulting in incorrect voltage requests to the processor".

It seems there was a firmware bug all along, where the chip asked for more power than it could handle. The initial BIOS fix lowered the strain by lowering performance, but hadn't addressed the underlying problem. The bug affects all Raptor Lake chips with a TDP of 65W or more, although it is only a real problem for chips that are pushed the hardest.

A microcode patch issued in August should fix things. However, the patch won't fix damaged processors that have been run beyond safe limits and physically degraded. Intel has offered to replace, under warranty, damaged chips. The standard Intel warranty is three years. Intel has extended this to five years for selected chips. If you have a retail chip, you should contact Intel. If your chip is part of a system then you must contact the manufacturer. There is to be no large-scale recall, no doubt because it's an expense Intel wants to avoid.

It's fair to say that Intel hasn't had a great time of late.

The replacement process requires you to send a picture of your chip for it to be verified, and your credit card details before it will ship a replacement. Your card is refunded after you've got your new chip, and Intel has received the damaged one. Needless to say, that is far from ideal. Unsurprisingly, this has attracted the attention of lawyers. A law firm is investigating the possibility of a class action lawsuit. There have also been concerns that some CPUs might harbor invisible damage. While Intel claims that currently stable chips will continue to be so after the patch, it's hard to be certain about that.

Raptor Lake crashing has been embarrassing for Intel. We don't know how many chips it'll have to replace, but it could get expensive, especially if everybody who has their Raptor Lake crash assumes the worst, and requests a replacement.

Intel has been here before with the Pentium FPU bug, which led to the first full-scale processor recall, and cost $475 million to fix back in 1994. This isn't the only problem Intel has—its financial results for the second quarter are in. The company has lost $1.6 billion, and doesn't expect the next set of numbers to be much better.

This has prompted cutbacks and restructuring. The bottom line is the loss of 15,000 jobs, and a $10 billion reduction in expenses by 2025. Oh, and the cancellation of Intel's Innovation event in September, and the suspension of its share dividend.

Intel's breakneck development cycle was always ambitious. It also missed the rapid and profitable development of AI hardware, despite being among the first to build neural network chips. Intel is a big outfit, so it will weather this for now. It still has about 80 percent overall x86 CPU market share, but it needs to keep to its release schedule. The imminent launches of the Lunar Lake mobile and next-gen Arrow Lake desktop CPU families need to go smoothly. **–CL**

> "We don't know how many chips it'll have to replace, but it could get expensive

© INTEL

NEW RTX 4070 VARIANT

NVIDIA'S GEFORCE RTX 4070 is a popular card in the Ada Lovelace 40 series. It's the cheapest of what you might call 'proper' gaming GPUs in the Ada range, though it's still not truly affordable.

The RTX 4070 was joined by a zippier Ti, Super, and Ti Super versions. Now, we are to get a new variation of the vanilla RTX 4070. The new card swaps out the GDDR6X memory used in all 4070 or higher series cards for regular GDDR6, as used in the 4060. As far as we know, everything else remains the same.

The change means 20Gb/s modules with a maximum bandwidth of 480 GB/s, rather than 21Gb/s and 504GB/s. Performance is expected to drop by a sliver. Rumor has it that a batch of Micron GDDR6X modules failed quality control, leaving supplies tight. The new card keeps the 4070 available, while GDDR6X modules are saved for the 4080 and 4090 cards. If you are in the market for a regular 4070, read the technical specifications carefully. Here's hoping the slower memory at least makes for lower pricing. **–CL**

ZEN 5 HITS 6GHZ

SOMEONE GOING BY 'Skatterbencher' has been overclocking AMD's new Zen 5, and pushed a Ryzen 7 9700X to 5.8GHz on single-core. This is impressive, as it has a Fmax (frequency maximum) limit of 5.55GHz. He also reached 6.0Ghz using AIO water cooling. Much of this was achieved via AMD's Precision Boost Overdrive 2 toolkit.

Meanwhile, the Zen 5 halo chip, the Ryzen 9 9950X, has received a lukewarm reception. The $649 16-core chip is no slouch, but not fast enough to match Intel's current finest, the Core i9-14900K.

Games, by some measures, are barely one percent faster. Multi-core productivity is decent, but for a top-dog chip, it failed to enthrall many early reviewers, and AMD's Zen 5 launch has fallen a little flat. See what we think of its sister chip, the Ryzen 7 9700X, on page 74. **–CL**

Tech Triumphs and Tragedies

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

TRIUMPHS

MAGNETIC GATES
A paper has suggested that magnetic tunnel junctions could be used as faster and more efficient transistors.

MASSIVE DRIVES
SanDisk has revealed the world's first 8TB SD card, and 16TB external SSD drive.

CONSOLE COLLECTION
A retro gaming enthusiast has connected 444 consoles to a single TV—a Guinness world record.

TRAGEDIES

ARIZONA FAB WOES
Production at TSMC's Arizona plant has slipped. Practices like 12-hour days have not gone down well.

DON'T TAKE AI SERIOUSLY
Microsoft updated its Service Agreement to remind that AI is not "to be used as substitutes for professional advice".

RAM TO GET EXPENSIVE
High-bandwidth memory demands mean DRAM production is scaling back.

ARROW LAKE GETS BENCHMARKED

Engineering samples aren't impressive

AFTER THE STOP-GAP that was Intel's 14th-gen chips, we await the real next generation of desktop processors: Arrow Lake-S, or Core Ultra 200. We get a tile design, Intel's new 20A process node, a new socket, backside power delivery, and more. Launch is rumored to be October 10, and chips a few weeks later. We expect 14 SKUs, and the specifications of 12 of these have appeared so far, divided into three groups with TDPs of 125W, 65W, and 35W.

The top specification we know of is the Intel Core Ultra 9 285K. This has eight performance cores and 16 efficiency cores, while the base clock on the P-cores is 3.7GHz, with a boost to 5.4GHz, or 5.7GHz on a single cores. The lowliest chip so far is a Core Ultra 5 255F, which still packs six P-cores with a maximum boost of 4.7Ghz.

The upper mid-range Core Ultra 7 265KF with eight P cores and 12 E cores manages a Geekbench 6 single-core score of 3,219. This compares well with the Core i9-14900K at 3,088, though does not imply a huge step in IPC for the new core design. The multi-core score of 19,433 is slightly slower than the i9-14700KF chip the 265KF replaces, and shares the same core count.

Back in the mid-range, the Core Ultra 5 245F has been tested using Geekbench 5, and looks more impressive. It managed 2,248 on single-core, and 18,354 on multi-core. This is a 125W 6+8 chip with a maximum boost of 5.2GHz. It manages to roughly match a Core i9-13900K, the top chip from two generations ago.

Power consumption for Arrow Lake is rumored to be about 100W lower than the equivalent Raptor Lake chips. Luckily for Intel, AMD's Zen 5-based Ryzen 9000 chips are themselves a bit disappointing, so Arrow Lake has a chance. But it has to be stable, given the train wreck that is Raptor Lake's recent reputation. **–CL**

GOOGLE IS A MONOPOLY

So says the Department of Justice, but now what?

GOOGLE'S RISE to search engine dominance has been rapid. It has even become a transitive verb, and been added to the dictionary. Any company with such a grip on a market faces legal troubles, but this latest one has landed hard.

The DoJ has ruled that Google has violated antitrust laws, effectively becoming a monopoly for internet search services through illegal practices, violating Act 2 of the Sherman Act. The ruling opens the way for legally enforced measures to restrict the monopoly. It was also hailed as "a historic win for the American people" by US Attorney General Merrick Garland. The White House press secretary said Americans "Deserve an internet that is free, fair, and open for competition".

What's next is unclear. Among the options are divesting parent company Alphabet of the Chrome browser. The ownership of Android is also under review, given how it allows Google a tight hold over any smartphone that isn't from Apple. Untangling such assets from their parent company won't be easy.

Another possible measure is the ending of payments made to Apple and Mozilla for Google to be the default browser search engine. These are not insubstantial sums—Apple currently gets $20 billion a year. Compelling Alphabet to share ad data is also under review. More practical measures include stripping away some of the advertising revenue.

As with all such cases, things move slowly. Google has enough influence and money, to fight vigorously and effectively, so the search engine isn't going anywhere in a hurry. But Alphabet has been judged to have broken the law, which has to have ramifications. However, they won't be quick, and it probably won't really be what Alphabet really deserves, either. **–CL**

Geekbench AI 1.0

PRIMATE LAB'S GEEKBENCH is one of the biggest benchmarking tools going. However, the rise of AI means there's a gap in our assessment of PC performance. Enter Geekbench AI 1.0: a cross-platform benchmark that uses real-world machine learning to provide some hard numbers.

It runs ten AI workloads, each with three data types, and gives single-precision, half-precision, and quantized scores (32-bit, 16-bit, and 8-bit). Tasks include object detection, image classification, and image upscaling. It tests CPUs, GPUs, and NPUs, and runs identical tests on all hardware.

The point of reference is a Lenovo ThinkStation P340, which runs a Core i7-10700 processor. This is rated at 1,500. Version 1.0 supports a number of AI frameworks (the basic AI structure used to create the models). These include OpenVINO, ONNX, Qualcomm QNN, Samsung ENN, and Core ML.

Geekbench AI has been in development since 2021 under the name Geekbench ML. This initial release is expected to be rapidly updated as AI systems evolve. Speed isn't everything, however—accuracy is a factor, and Geekbench added the new accuracy measurements for more holistic scores.

Putting a hard number to PC performance isn't easy, as Primate Labs' founder, John Poole, posted: "AI takes this problem and turns up the complexity dial to 11". Geekbench won't tell us which AI implementation gives the best summaries of text, or the most artistic images, but does give us a point of reference for comparisons. If you want to see how much AI horsepower you're packing, Geekbench AI is available for download now, and the results list is up and running. **–CL**

Cheap AMD V-Cache

Often, the first we hear of a chip is in a product database, or documents seeking regulatory approval before official announcement. This is where we are with the AMD Ryzen 5 5500X3D. It's not a guarantee of existence, but it has appeared in paperwork for the Eurasian Economic Commission.

The new arrival is a Zen 3 chip for AM4 socket boards, proving AMD's commitment to older standards. It has six cores, 12 threads, and a base clock of 3.0GHz, plus a top boost speed of 4.0GHz. It also has AMD's secret weapon for gaming goodness, V-Cache, which layers a fat 80MB of L3 over the cores for a total of 96MB. It has a slightly lower clock than the standard Ryzen 5 5500, but that extra L3 should compensate. The Ryzen 7 5700X3D launched at $249, and is now available for $200-229. The 5500X3D should be cheaper, making it a budget gaming rocket. **–CL**

China's H100 AI Rival

Huawei is testing its new AI chip, the Ascend 910C, pending full launch. The 910B sold well, as Chinese firms had little choice—thanks to trade sanctions, they can't (legally) get Nvidia's storming H100 chips. The fastest AI GPU the US government allows for export to China is the Nvidia H20. This has only 41 percent of the H100's core count, but China is expected to buy $12 billion of them.

The Ascend 910C is Huawei's answer, and early reports claim it is 'about' as fast as the H100. A bold claim, perhaps, but adversity is the mother of invention, and US sanctions have been in place since 2018, providing plenty of incentives to accelerate chip development.

The 910B sold about 550,000 units in China, despite some technical defects, against the H20's million plus. Eventually, China will be able to produce better chips than it can import under restrictions. That's surely just a matter of time. **–CL**

Jarred Walton

TECH TALK

Intel issues fix for Raptor Lake degradation

EARLIER THIS YEAR, I wrote about difficulties I was having with a Core i9-13900K processor (see MPC230 Tech Talk). Little did we realize that we were only seeing the tip of the iceberg. While most complaints have involved the unlocked Core i9 Raptor Lake CPUs, it appears the instability problems build up and potentially impact many Raptor Lake—13th and 14th Gen Core—CPUs, with Intel identifying 22 different desktop parts.

Everything from the Core i5-13600K and above rated for more than 65W base power, including i7 and i9 65W chips, is affected. That means T-class parts should be fine, as well as vanilla models like the Core i5-14600, but popular unlocked K-class processors are impacted. There are also reports of problems extending to server and laptop parts.

Intel has a microcode fix that has begun rolling out via BIOS updates. It supposedly targets the root cause, but if your CPU has already showed problems, there's no fix other than replacing it. Intel isn't issuing a recall, but I can smell the class action lawsuits. However, it has extended warranty support on the affected CPUs by two years. The official statement reads, in part: "Intel's current analysis finds there is a significant increase to the minimum operating voltage (Vmin) across multiple cores on affected processors due to elevated voltages. Elevated voltage events can accumulate over time and contribute to the increase in Vmin for the processor. The latest microcode update (0x129) will limit voltage requests above 1.55V as a preventative mitigation for processors not experiencing instability symptoms."

Our understanding is that this goes back to the integrated voltage regulator (IVR) in Raptor Lake CPUs. Modern processors are designed with some 'intelligence' when it comes to voltages and the IVR monitors functioning of the chip, and can increase the delivered voltage. Two chips of the same CPU model can have small low-level differences, so having the hardware detect what's required and adjust normally works. But Raptor Lake may have pushed clock speeds and voltages too far, and when the IVR sends higher than expected voltages, the result appears to be a cascade effect.

Raptor Lake processors from the Core i5-13600K up to the Core i9-14900K are potentially impacted by voltage spikes.

Over time, the effect can cause the IVR to slowly increase the baseline voltage. But the spikes continue, further degrading the CPU structures, then the IVR has to keep edging upward on voltages. Eventually, it reaches the critical limit, at which point you basically have a paperweight CPU.

That seems to be what happened to my 13900K. Eventually, the chip reached the point where I couldn't install Nvidia's GPU drivers, and other applications began misbehaving. A clean install of Windows 11 crashed before I could get things up and running. At that point, I had to RMA the CPU.

The new chip has been working fine. Intel also says that the microcode updates shouldn't affect performance—this is about limiting voltages, and high transients weren't helping. Future CPUs should also be unaffected—these issues appear to be specific to the Raptor Lake generation. But the damage to Intel's reputation will be harder to overcome.

I can't help but think some of this goes back to motherboards pushing the limits a bit too far. One PC vendor, Puget Systems, even says that it hasn't noticed any clear indication of problems with the Raptor Lake systems it shipped, but it also notes that it applied more limited power profiles than other companies. Don't be surprised if future Intel chips come with more strongly enforced power and voltage limits to prevent similar problems from occurring.

> "Intel isn't issuing a recall, but I can smell the class action lawsuits

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

THE LIST

THE BEST MOUSE PADS FOR GAMING

THE BEST MOUSE PADS offer slick movement and awesome aesthetics, no matter what surface you prefer. At the most basic level, you want something that offers a smooth and precise surface for your mouse, which thankfully can be bought for not much more than a couple of coffees. Beyond that, you might want features like USB passthrough, RGB lighting, and maybe even less of a mouse pad and more of a deskpad. Here are our favorites.

5 NOVELKEYS DESKPAD

If you're the sort of person who feels constrained by a traditionally sized mousepad, then why not try a deskpad? NovelKeys does a lot of custom designs based on themes like *Ghostbusters* or *Star Wars*, but they feel good to use and last a long time. Most mouse pads come in at an accommodating 900 x 400 x 4mm size, with a rubber bottom, cloth top, and outer edge stitching to match.
$13-75 www.novelkeys.xyz

4 LOGITECH G440

Nothing compares to a solid, hard surface mousepad when it comes to pure speed and low tactile resistance, and thanks to Logitech, you don't have to pay the earth to get those things. There are lots of different hard surfaces available, but we like the simplicity and ultra-smooth surface of the Logitech G440. The polyethylene surface features microtextures that offer the perfect amount of resistance, providing just the right level of feedback without any excess noise.
$20 www.logitechg.com

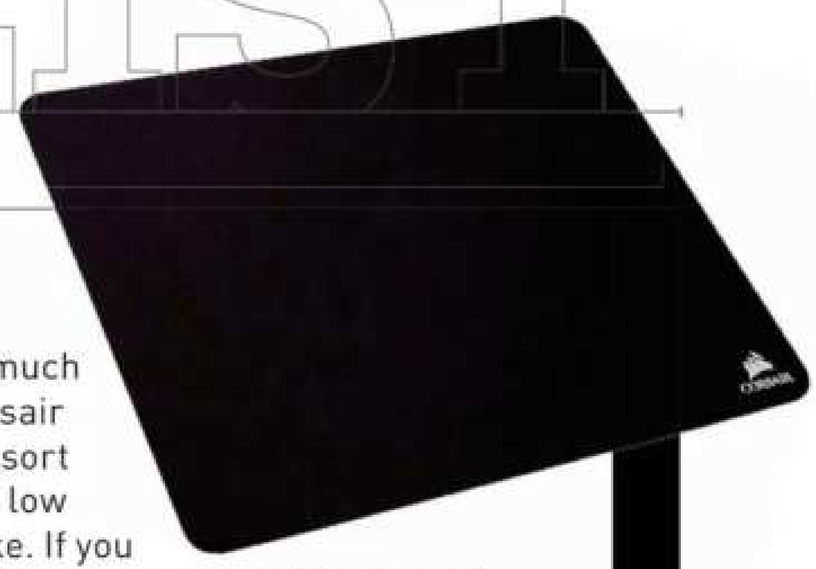

3 CORSAIR MM100

There's not much to say about the Corsair MM100, but it's that sort of simplicity, and its low price tag, that we like. If you want a nondescript black rectangle with minimal branding that your mouse will glide across, then this is your pad. It's smooth on top and rubberized on the bottom. The only downside is there's no stitching on the edges, so repeated scuffs can cause some frays. **$10 www.corsair.com**

2 RAZER STRIDER

Think of this as a hybrid mouse pad, offering the flex and soft finish of a fabric mat with the stick-to-itiveness and glide of a more rigid model. The reason we love it is the anti-slip rubber base that clings to your desk, with zero chance of lateral movement. For a no-fuss and well-made pad, look no further than this.
$50 www.razer.com

1 RAZER FIREFLY V2 PRO

There's no finer a mouse surface than the Razer Firefly V2 Pro, a hard-top mat that also comes with some very alluring RGB lighting. Of course, you don't need the latter, but you'll love the way your mouse moves as if it's a hovercraft. It's also easy to wipe down and clean with a damp cloth. This is the best mouse pad we've ever used.
$100 www.razer.com

© NOVELKEYS, LOGITECHG, CORSAIR, RAZER

Jeremy Laird

TRADE CHAT

AMD's turn to drop the ball?

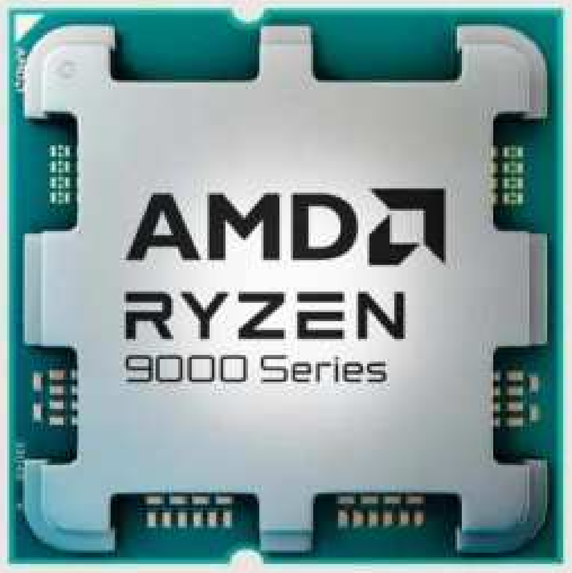

Just when AMD was sprinting into Intel territory, it has fumbled with the new Ryzen 9000 CPUs.

WITH INTEL'S RAPTOR LAKE CPUs falling over, the company firing around 15,000 employees, and cancelling its 2024 innovation event, AMD must have been enjoying the view—until its new Ryzen 9000 desktop CPUs rolled out. So, is AMD's CPU a minor stumble or game-changing fumble?

When it comes to CPUs, as opposed to GPUs, AMD just keeps on rolling. The first Ryzen-branded CPUs launched way back in summer 2017. Admittedly, the Ryzen 1000 family and its Zen architecture didn't immediately put AMD in a leadership position, but it was far more competitive than the chip it replaced.

From there, AMD iterated, improved, and generally delivered. Each new Ryzen generation was an unambiguous step up from its predecessor. Okay, some were more impressive than others, but there were no significant missteps.

That's impressive, but it arguably comes to an end with the new Ryzen 9000 generation. Certainly, Ryzen 9000's 'new' Zen 5 architecture is disappointing by pretty much every measure. When it comes to IPC or instructions per clock, it's essentially a single-digit percentage improvement.

Sure, any improvement is welcome, but if this was typically the pace of improvement, imagine how many generations, and more to the point, years, you'd need for the kind of unambiguous upgrade that you can really feel in the real world.

Speaking of real-world performance, some games actually show a regression in performance versus Ryzen 7000 CPUs—not by enough that you'd notice, and only in a few titles—but it's hardly indicative of a really successful new CPU design.

There's talk of issues around internal latency impacting performance, but whatever the reason, there's no getting round the fact that Ryzen 9000 is a bit of a disappointment. That's doubly true when you consider that AMD is sticking with the same core counts as before.

> Certainly, Ryzen 9000's 'new' Zen 5 architecture is disappointing by pretty much every measure.

Put it this way, in a blind test, there's no chance you'd be able to feel the difference of an 'upgrade' to Ryzen 9000 in games, or any other application. In mitigation, the new chips can be quite a bit more efficient in some scenarios. That's going to be a tangible benefit for server applications, and it will probably be enough to maintain AMD's steady erosion of Intel's market share.

Back in the PC, it's not worth all that much. What's more, the Ryzen 9000 doesn't seem to be fully baked at launch. As I write these words, it's not officially confirmed, but it looks like AMD is going to release a microcode update that may actually increase the TDP of certain models in the new 9000 range.

If that helps put some air between the new processors and last-gen chips, it will be welcome. But honestly, customers shouldn't have to worry about updating their CPUs weeks after buying them. Likewise, AMD was under no pressure to roll these new CPUs out, and it must have known how they performed and that little to no improvement in many benchmarks was hardly going to be met with widespread applause.

It all feels a little like the AMD of old; the one that struggled to get almost any new CPU generation out the door without tripping itself up in the process. So far, it's not really a disaster, and I have a hunch that AMD is actually going to get away with this one. For starters, you'd still buy AMD rather than rolling the dice on an Intel Raptor Lake CPU that may or may not truly be fixed by Intel's own microcode patches.

The real test of whether Ryzen 9000 and Zen are a flop will come in a month or so when Intel's own new CPU, known as Arrow Lake, is due to appear. But all the indications are that Arrow Lake, like Ryzen 9000, is a disappointing step over existing CPUs. If so, Ryzen 9000 will be competitive in performance terms, and won't suffer the reliability concerns that will lumber Intel CPUs for a while to come.

And yet, Ryzen 9000 is a pity for AMD. To mix my sporting metaphors, just as it has Intel on the ropes, AMD has fumbled the ball. Maybe it won't matter. But maybe, just maybe, it will give Intel that little bit of breathing space it needs to get back in the game.

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

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> Music streaming
> Switching to Linux
> Duplicate removal

Access web vault

I read with interest your advice on how to use your WireGuard VPN to make it possible to take sensitive services 'offline' by making them accessible only through your network. My only disappointment was that it prevents me from accessing my Vaultwarden web vault remotely due to the issue with https. Is there really no workaround for this?

—Sean F Maxfield

THE DOCTOR RESPONDS: The Doc has done some extra digging and found one workaround you can use involving a domain name you own. The Doc assumes that you've followed our previous tutorials on configuring subdomain redirection using Nginx Proxy Manager, and that in this instance you have a subdomain (such as bw.domain.com) set up with an SSL certificate already attached to it.

In the short term, all you need to do is instruct your domain provider to point your subdomain to a local IP address instead of having it redirect to your public IP address. Log into your domain provider's control panel and locate the entry you created for bw.domain.com. Delete this and then create a new 'A record' (not 'A+ Dynamic DNS Record') for bw.domain.com, pointing to your proxy server's local IP address (such as 192.168.0.1). Leave your existing proxy server redirection in place—don't disable or delete it.

This means any device attempting to connect to bw.domain.com is redirected to 192.168.0.1 on its own local network, which means the domain only works when anyone is connected at your home or through your Wireguard VPN server.

This leaves one problem: When the time comes for Let's Encrypt to renew bw.domain.com's SSL certificate, the HTTP challenge will fail because it can't access your Vaultwarden server via your domain provider. The band-aid approach would be to temporarily redirect bw.domain.com back to your public IP address from your domain provider's control panel. Run the SSL renewal again from inside Nginx Proxy Manager—after it succeeds, redirect bw.domain.com back to 192.168.0.1 until the SSL certificate next needs renewing in 90 days' time.

If you're feeling brave, there's a potentially permanent fix you can try. This involves creating a wildcard SSL certificate (*.domain.com) that can be linked to all your subdomains, including bw.domain.com. The problem with this approach is that you need to configure your setup to allow a DNS challenge instead of an HTTP one, which is more complicated. Visit https://letsencrypt.org/docs/challenge-types/ for more details of this approach.

If your domain/DNS provider is listed, obtain the API token following your provider's instructions, then log back into Nginx Proxy Manager, switch to the 'SSL Certificates' section, and click 'Add SSL Certificate'. Type '*.domain.com' (replacing domain.com with your own domain) and flick the 'Use a DNS Challenge' switch on. Now select your domain provider from the drop-down menu and input details of the required API token as shown. All being well, on clicking 'Save' you're granted a wildcard certificate you can use not just for Vaultwarden but any other services you'd like to take 'offline' in a similar manner going forward.

Cheap music streamer

Can you suggest a cheap solution for streaming music on my network? I'd like to cannibalize existing components like my hi-fi if possible. **—James D Gusman**

THE DOCTOR RESPONDS: The Doc is a long-time fan (and user) of the open-source Logitech Media Streaming platform. It's recently been renamed Lyrion Music Server (https://lyrion.org) with a user interface refresh, but the underlying tech remains the same. It's a three-tier structure: server, controller, player.

You can run the server on any hardware, including your main Windows PC or self-built NAS server running Linux (either natively or through Docker/Podman). Full instructions and downloads can be found on the website under 'Getting Started', but if you're a Podman user, the following should get you up and running (adjust the '-v' lines to point to your personal folders as needed):

↘ submit your questions to: doctor@maximumpc.com

```
podman run -it -d \
    --name=lms \
    --net=host \
    -v "/home/yourname/
container-data/logitech/
config":"/config":rw \
    -v "/home/yourname/
Music":"/music":ro \
    -v "/home/nickdanp/
container-data/logitech/
Playlists":"/playlist":rw \
    -v "/etc/localtime":"/etc/
localtime":ro \
    -v "/etc/timezone":"/etc/
timezone":ro \
    docker.io/lmscommunity/
lyrionmusicserver:dev
```

You then access the server through a controller, which can be any web browser via http://192.168.x.y:9090 (substitute x.y with your server's IP address), a mobile app or dedicated piece of hardware, such as a Raspberry Pi with a touchscreen.

The controller can be used to direct output to any player on your network, from your phone or tablet to Bluetooth audio receivers and Logitech Squeezebox devices you might still have gathering dust somewhere. You can also build a dedicated music player to output through a hi-fi—we recommend pairing a Raspberry Pi of any vintage with a suitable audio DAC. For example, you can source a Raspberry Pi Zero 2 W with pre-soldered headers, HiFiBerry DAC+ ZERO with phono connectors and essential extras (case, microSD card, and microUSB power supply) from www.pishop.us for under $75 plus shipping. Pair this with PiCorePlayer (https://picoreplayer.org) and you have yourself an audiophile-friendly player.

Ditching Windows

I'm fed up to the back teeth of Microsoft's ever-increasing assault on privacy. I'd like to switch to Linux full time but am aware many key apps don't come with Linux builds. What would my options be if I wanted to force this switch?

—Ann Reyes

THE DOCTOR RESPONDS: We feel your pain, Ann. Without serious pushback, Microsoft will simply continue to poke its nose further into Windows users' activities to monetize them. Linux is the ultimate antidote to this, but the problem is finding a practical workaround for apps that won't run in Linux.

Attempts to emulate Windows directly through a genuine open-source OS have resulted in ReactOS (https://reactos.org). Sadly, this has never developed beyond the alpha stage and while you might like to play with it in a virtual environment, it's not suitable for use on a daily basis.

This leaves you with macOS or Linux. The former requires an expensive outlay on Apple gear, but you'll find more native support for popular apps like Office. Linux is completely free, but your first job will be to choose a suitably accessible flavor—we recommend Ubuntu or an Ubuntu-derivative such as Linux Mint (www.linuxmint.com) or Zorin OS (https://zorin.com) for Windows switchers.

From here you then see if you can source usable open-source alternatives to any core apps that aren't supported in Linux, such as LibreOffice (www.libreoffice.org) or GIMP (www.gimp.org). Where none can be found, try Wine (www.winehq.org) or one of its more user-friendly front ends, such as CrossOver ($60 after free trial, www.codeweavers.com) to see if your app can be run using a Windows compatibility layer.

You'll typically have more success with older versions of apps, such as pre-2016 flavors of Microsoft Office as well as the Windows build of Steam, which opens up access to many Windows-only games found on the Steam platform. Use CrossOver's compatibility database (www.codeweavers.com/compatibility/) to see what might be possible.

Your final option is to set up a virtual machine inside Linux using VirtualBox (www.virtualbox.org), inside which you'd set up a Windows installation to continue running those apps that can't be emulated or replaced.

Tenacity vs. Audacity

In Nate Drake's otherwise excellent article on Audacity 3.5, he talks about the privacy issues surrounding the program and mentions Tenacity as an alternative. However, he didn't specify the key differences between the two. Does Tenacity follow Audacity releases, or has it forged its own separate path?

—Frank S Lauber

THE DOCTOR RESPONDS: Tenacity's own history is complex as it was one of several Audacity forks made after the Muse Group acquired Audacity in April 2021. Within a year, most of these—including Tenacity itself—had virtually ceased to function. Eventually, two other projects merged with Tenacity to resume development. However, at time of writing, the project hasn't been updated in over nine months, following the minor 1.3.3 release.

According to Tenacity's own release notes, its first major stable release (1.3) compared itself to Audacity 3.1, released in October 2021. Since then, Audacity has rolled out five major point releases, each containing new features. For example, version 3.6 added support for master effects as well as a new compressor and limiter. In other words, despite the controversy over privacy, Audacity has evolved significantly in positive ways since the Muse Group first took control.

Visit https://support.audacityteam.org/additional-resources/changelog/older-versions and make a note of the new features added since version 3.1 to help determine which app you wish to use. As Nate states, you can always configure your firewall to block Audacity's access to the Internet if you don't trust it, although you lose access to its new cloud saving feature if you do.

Consolidate folders

I'm looking for a tool to go through folders and strip out duplicate and visually similar photos. Most of the tools I've seen come with a price tag attached – can you recommend any freebies?

—Mary Martelli

THE DOCTOR RESPONDS: There are several free duplicate image finder tools, but one of the simplest is Awesome Duplicate Photo Finder (www.duplicate-finder.com). Select one or more folders, click 'Start Search', and wait for the program to run through the folder—you're shown how many photos have been found and scanned, and the number of visually similar images that have been found.

Once the scan is complete, a list is shown in the bottom pane. Check the Similarity column, which gives each match a similarity score—the higher the score, the closer the match. Click the column header twice to sort the results with the closest matches shown first.

Select the top entry and you'll see a preview of each image, complete with image stats (such as resolution, size, and file type). You'll also see options to move the file elsewhere, delete it, or browse the parent folder should you need context about this particular copy of the photo.

Other tools offer more options, such as setting 'protected' folders, but add complexity—our favorite alternative is AntiDupl, which is also free and open source (https://github.com/ermig1979/AntiDupl).

Phantom of the Opera

CPUs Unmasked

Zak Storey explores why Zen 4 might be the best CPU generation to date

LET'S FACE IT, this industry isn't exactly one that has avoided controversy over the years. The PC enthusiast landscape has always been riddled with, let's call them, issues. In fact, right now, we're witnessing one of those very moments rising up again, and most notably, that comes in the form of Intel's hiccups (putting it politely) with its 14th and 13th generation of processors. The TL;DR is that the manufacturing process in the 13th generation chips potentially led to oxidization across certain transistor gates, leading to CPU malfunction, and the 14th gen's micro-code, under extreme load and voltages would actively cause the chip to deteriorate with no hope of recovery once the damage had been done. All of this only came to light after extensive investigations by some seriously admirable media and B2B partners (thank you, Gamer's Nexus).

Intel has aggressively pursued micro-code tweaks since the first acknowledgement of its 14th-gen stability faux pas, and as such, the chips themselves 'should' be more stable, and less likely to commit seppuku on themselves under heavy loads. That said, it's not quite done yet, and we're still expecting a BIOS tweak in the coming weeks. If your 14th-gen chip is having problems, please get in touch with Intel's customer support directly, and they should be able to help you out.

As we've already had one major round of BIOS updates, forcing motherboards into an 'Intel recommended' voltage config, and with micro-code updates landing in just two weeks time, we've decided to do a two-parter this month: one build comparing Intel's current 14th-gen state against Ryzen 7th gen under these new parameters, and one build next month, still comparing the same 14th-gen build, but now with the new micro-code changes, with Ryzen 9000 as our feature build as well.

To round out our performance analysis, we're also taking this opportunity to launch the Maximum PC System Benchmarking Scheme 2.34. We've massively overhauled our testing for PC builds, and moving forward, we'll be going from 10 to 50 data points across 20 different tests. Each system tested will include everything from storage results, in-depth CPU analysis, AI inference benchmarking, temperature and power draw figures, along with 1080p and 4K gaming results, to give all of us a far better picture of just how well these systems perform. So with all that said and done, let's jump right into what parts we've picked for (possibly) our final ever Ryzen 7th-generation PC.

SPECIFICATIONS

Part		Price
CPU	AMD Ryzen 9 7900X	$359
Mobo	NZXT N7 B650E Motherboard	$290
CPU Cooler	NZXT Kraken Elite 360	$300
RAM	32GB (2x16GB) Kingston Fury Beast DDR5 @ 6,000 MT/s	$110
SSD 1	1TB Samsung 990 Pro M.2 PCIe 4.0 SSD	$110
SSD 2	4TB Western Digital SN5000 M.2 PCIe 4.0 SSD	$280
GPU	AMD Radeon RX 7800 XT	$510
Case	NZXT H7 Flow RGB 2024	$150
PSU	1000W NZXT C1000 80+ Gold	$180
Fans	2x NZXT F360 RGB Core Frame	$140
Total Price		$2,429
Core Price*		$1,379

*Core Price refers to the key components generating performance (CPU, GPU, Mobo, OS SSD 1, RAM), not accessories.

PRICES CORRECT AT THE TIME OF PUBLISHING

https://content.jwplatform.com/videos/K5lw0fGb-u2lN49He.mp4
Please type this URL into your browser if the link is broken

STEP-BY-STEP GUIDE
BUILD THIS PC!
PG. 21

GREAT EXPECTATIONS

CPU

AMD RYZEN 9 7900X

WWW.AMD.COM

What we've seen from early looks at the 9700X and 9600X doesn't fill us with much hope in terms of massive IPC jumps. That makes the 7th gen an increasingly interesting value proposition right now. For $359, you can pick up the 7900X, complete with 12 cores, 24 threads, and a maximum boost clock of 5.6GHz.

It's also remarkably cool, with a relatively low power draw, along with some impressive efficiency.

If gaming or single-core prowess is more your thing, you may opt for the 7900X3D with its expanded V-Cache (going from 64MB of L3 to 128MB), but that is $35 more than this little number. • **$359**

RAM

32GB (2X16GB) KINGSTON FURY BEAST DDR5 @ 6,000 MT/S

WWW.KINGSTON.COM

Ryzen's Infinity Fabric is AMD's name for its interconnect that surrounds each CPU's core complexes, joining them together and transferring data. In the early days, when Ryzen had a larger number of CCDs, that interconnect improved its performance and reduced the latency between those CCDs, allowing them to communicate faster.

Today, although our Ryzen processors have fewer CCDs, they still have them. The 7900X has two. Similarly to previous gens, that Infinity Fabric is tied to the speed of the memory controller. Boost the controller's speed, and wham, faster CPU.

Once Ryzen 7th gen hits above 6,000 MT/s, the memory controller frequency decouples from RAM speed, and we're well into the realms of diminishing returns. That's why we've gone back to Kingston's super-budget 6,000 MT/s kit. 32GB is plenty for what we need, and that 6,000 MT/s speed under EXPO makes this a surefire pick to maximize our Ryzen potential. • $110

CPU COOLER

NZXT KRAKEN ELITE 360 RGB

WWW.NZXT.COM

We have another NZXT pick this month, this time in the form of the Elite 360 RGB. Let's be up front—we actually requested a slightly cheaper model from NZXT, but unfortunately, press samples were limited at the time (not helped by an aggressive deadline), so for $300, it's a lot for a cooler of this caliber.

It's still great, of course. You get that LCD display, decent RGB fans, and plenty of capacity, but if you want to save a bit of cash, it may be smarter to grab the standard Kraken 360 ($180) or Elite 360 non RGB ($267), both of which shave off some serious cash, and you still get a fancy display in the middle of them for all the stats or GIFs you can muster.

That's not to say the Elite 360 RGB is bad, of course (it isn't, and we're more than happy to feature it here), as it still packs in an awesome pump, and a slightly slicker display than the other units, along with those RGB fans as well. • **$300**

MOTHERBOARD

NZXT N7 B650E

WWW.NZXT.COM

Looking for a solid AMD motherboard with a decent feature set and clean aesthetic? This might just be the ticket. NZXT's N series of motherboards have always impressed us here at *Maximum PC* (when we're not getting the PCIe configurations mixed up), and the B650E is no exception.

With a beautiful, crisp, armor-plated finish, available in black or white, it also stacks in support for DDR5 up to 7,200 MT/s, along with PCIe 5.0 support for the GPU and the top-most M.2 PCIe SSD slot (we triple checked). You also get 2.5G Ethernet, WiFi 6E, and Bluetooth 5.2, along with an arsenal of USB 3.2 Type C and 2 support, and a not unreasonable rear I/O, complete with 8 x USB Type A, 1 USB Type C, 5.1 Audio Out with S/PDIF, Clear CMOS and BIOS Flashback buttons, and an HDMI port. Honestly, how good is this B650E, and why is it better equipped than the Z790 board? • **$290**

SSD 1

1TB SAMSUNG 990 PRO M.2 PCIE 4.0 SSD

WWW.SAMSUNG.COM

The 1TB Samsung 990 Pro is still a bit of a beast, even compared to PCIe 5.0 standards. We'll talk a bit more about this later on, and just how close it gets to topping Crucial's T700, but the long and the short of it is that its Random 4K performance is outstanding. Combine that with some exceptional software through Samsung Magician and a decent(ish) warranty, along with 1.5 million hours before failure rate, and it's a solid pick, certainly at that price. You can even get it with a heatsink if you want, although it's more than comfortable operating all the way up to 70 C, which would be seriously toasty for a 4.0 drive under a thick, integrated heatsink. • $110

STEP-BY-STEP GUIDE BUILD THIS PC! **PG. 21**

SSD 2

4TB WESTERN DIGITAL SN5000 M.2 PCIE 4.0 SSD

WWW.WESTERNDIGITAL.COM

At the other end of the spectrum, we have WD's latest SN5000. Still under that PCIe 4.0 standard, and with a five-year limited warranty, unlike the Sammy drive this thing can skyrocket up to 85 C, and has the added bonus of a 1200 TBW before failure rate, too.

The big element? The pricing, WD currently have these on sale for just $280—that's around $70 a TB on a single, highly reliable drive, with plenty of sequential and random 4K performance to back it all up. It's not the quickest drive, not the latest specced, but it costs around the same as a 2TB PCIe 5.0 drive, only with twice the capacity. That's nothing to snub your nose at. • $280

GPU

AMD RADEON RX 7800 XT

WWW.AMD.COM

The Radeon RX 7800 XT represents surprisingly solid value. AMD has come a long way with its FSR tech, and the upscaling is potent in a number of titles, so Nvidia and Intel cards can take advantage of it. It's not a patch on DLSS or Nvidia's AI Frame Generation tech, but can lead to some serious frame-rate improvements without sacrificing much quality.

Core spec is solid, clock-speeds top out at 2.43GHz, it's got 60 compute units, comprised of 3,840 stream processors, 96 ROPs, 60 ray accelerators, and 120 AI accelerators. This is bundled with 16GB of GDDR6 on a 256-bit bus, driven by twin 8-pin PCIe power and the power from the slot. If 1440p or 4K gaming is your aim, this straddles that edge. • **$510**

FANS

2X NZXT F360 RGB CORE FRAME

WWW.NZXT.COM

The final piece of the NZXT puzzle are these F360 RGB Core Frame fans. To quote our photographer, "Why has no one done this sooner?" and he has a point. NZXT now sells a set of 'framed' fans. You can get them in either 240, 280, or 360 configurations, in black or white. You get a single cable connecting all of them, and a splitter to run off an aRGB header and a fan header, and you're set. Alternatively, if you've got an NZXT header on your motherboard, or an NZXT fan controller, that'll fit the bill, too, and allow you full individual control on all of the addressable RGBs.

Installation is a breeze, with just four screws—one on each corner—and stats are pretty solid, too, with 3.3 mmH2O per fan combined with 75.12 CFM per fan at full tilt. Noise tops out at 30 dBA. • $140

PSU

1000W NZXT C1000 80+ GOLD

WWW.NZXT.COM

This is definitely an overkill pick. We actually recommend an 850W PSU if you are building something similar. Still, 1000W isn't a bad choice, certainly if you intend to upgrade the GPU. There's no telling how high the power draw will be on Nvidia's next set of GPUs, so something with more head room might be a good choice.

From the wall, PC Part Picker reckons the maximum our system will pull is around 597W, and that's in a worst-case scenario, where all power draws are maxed out. It doesn't quite hit that in reality, but nonetheless, having the head room is nice. Otherwise, there are plenty of connectors, a zero dBA mode, and a 10-year warranty. Sadly, NZXT don't offer any custom cable kits, so if you want something a little slicker, you'll need to head to CableMod. However, even a non-pro set will put you back $120 before shipping if you go down that route.• **$180**

CASE

NZXT H7 FLOW RGB 2024

WWW.NZXT.COM

We waxed lyrical about the H7 Flow last issue, but it's such a slick chassis, we had to give it a bit of a showcase. This new iteration of the legendary H7 line comes in two colorways, and in two variants as well. You can have it in black or white, and your choice of RGB or non-RGB. If you go the RGB route like we (accidentally) have, you'll be graced with one of NZXT's latest Core Fan Frames as well in the front. More on that in a bit.

Either way, for $150, this thing absolutely slaps when it comes to cooling. You've got space for ten fans total, plus a 360mm rad in the roof and a 420mm in the front, combined with an absolute arsenal of cable management in the rear. On top of that, airflow is impressive across the board, and the general build quality is top notch, too.

The only downside? NZXT does technically sell a vertical GPU bracket for its H series cases, but it isn't compatible with the latest edition of the H7 Flow.
• $150

SIMPLICITY IS THE ULTIMATE SOPHISTICATION

STEP-BY-STEP GUIDE

LENGTH OF TIME: 1-2 hours
DIFFICULTY: Easy

OKAY, HANDS UP, this is perhaps one of the easiest builds we've pieced together in the last few months, maybe even in the last year (bar perhaps the integrated GPU build). There's no crazy new cable standards, no insane RGB cabling to deal with, and no wild and wacky cases to contend with. It's just a plain old standard PC, complete with some juicy storage and a super-simple build process.

If you're new to PC building, or are nervous about building a second rig, given how challenging your first one has been, this system here is a fantastic way of diving back into your hobby, and realizing how easy it can be—if the parts and hardware are designed well enough, that is.

This build honestly feels like a testament to how well NZXT has fine-tuned the H7 Flow and its corresponding components. There's little to do or worry about, and that's no bad thing.

If anything, it gives us more time to go over the finer details, and dive a little deeper into all of the benchmarking and CPU shenanigans that are currently ongoing. With that said, let's get going.

THE END AND THE BEGINNING

It's a thing we never really go into detail on, but at the start of every system build, the *Maximum PC* build team typically do some initial prep long before we even set foot taking hardware out of boxes. Once the constant flurry of emails back and forth between companies has subsided, and enough hardware has shipped, first and foremost our job is to get everything together and prepare our build space. That's a fine thing to consider when you go about building your own system.

Grab all of your hardware, and make sure it's all readily available. Then, you're going to need the tools necessary to piece everything together. A good set of screwdrivers or a multi-tool set will go along way. A pair of scissors or cable snips, along with some cable ties, will also come in handy, along with a good-sized table, desk, or some easily accessible workspace for you to build your system on.

Anti-static prep? Outside of some basic steps, we wouldn't worry about it much. Make sure you're not wearing wool, or woolly socks and sliding along carpet, and you should be fine. If you're really nervous, you can grab a metal pole, a radiator, or anything that's grounded to give yourself a good discharge, and you'll be good to go.

After that, it's time to get to work. The first thing, of course, is going to be getting that chassis out of the box (ironically where you're most likely to get a zap or two). Once out, it's time to perform that *Maximum PC* step one staple: the chassis teardown. We're going to be removing every panel we can, and anything from the case that we're not going to use **(Step 1)**. That can be glass windows, rear side panels, dust filters, hard-drive caddies, you name it—anything that might get in the way of your cable management or hardware.

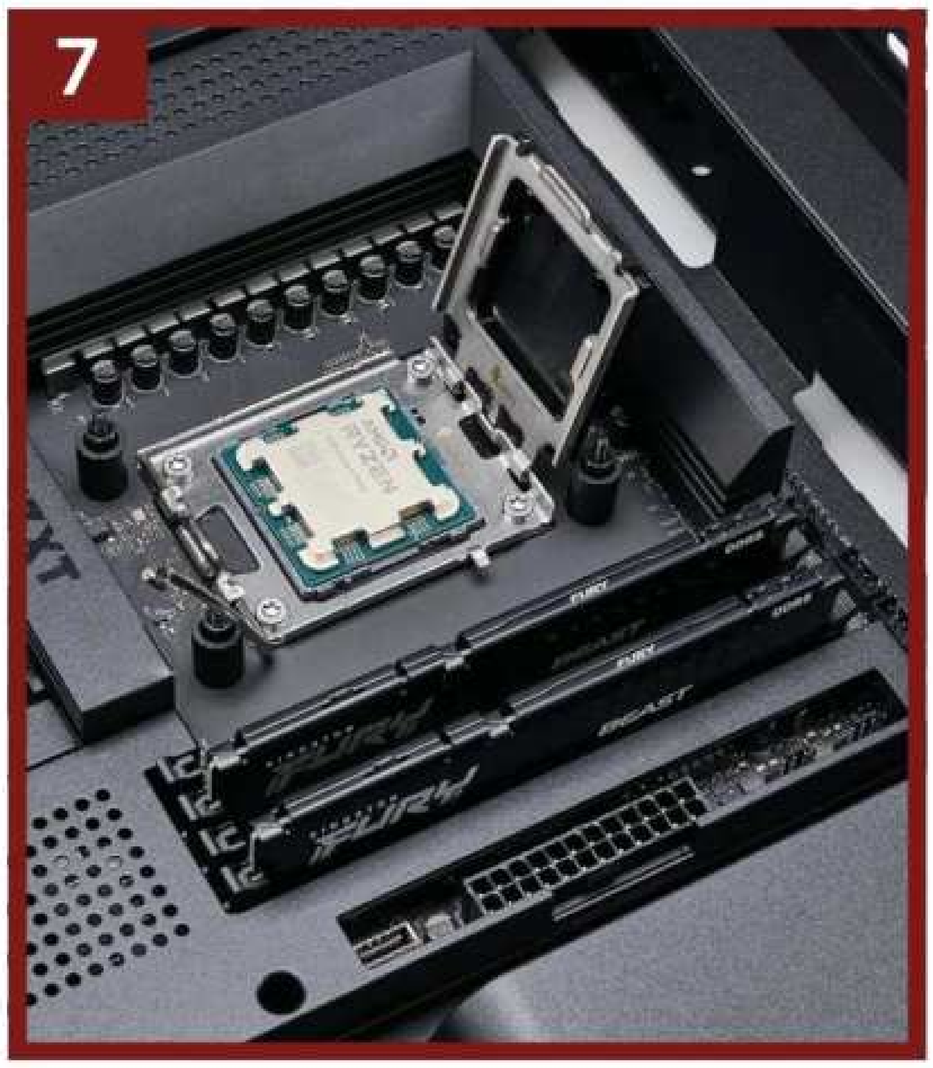

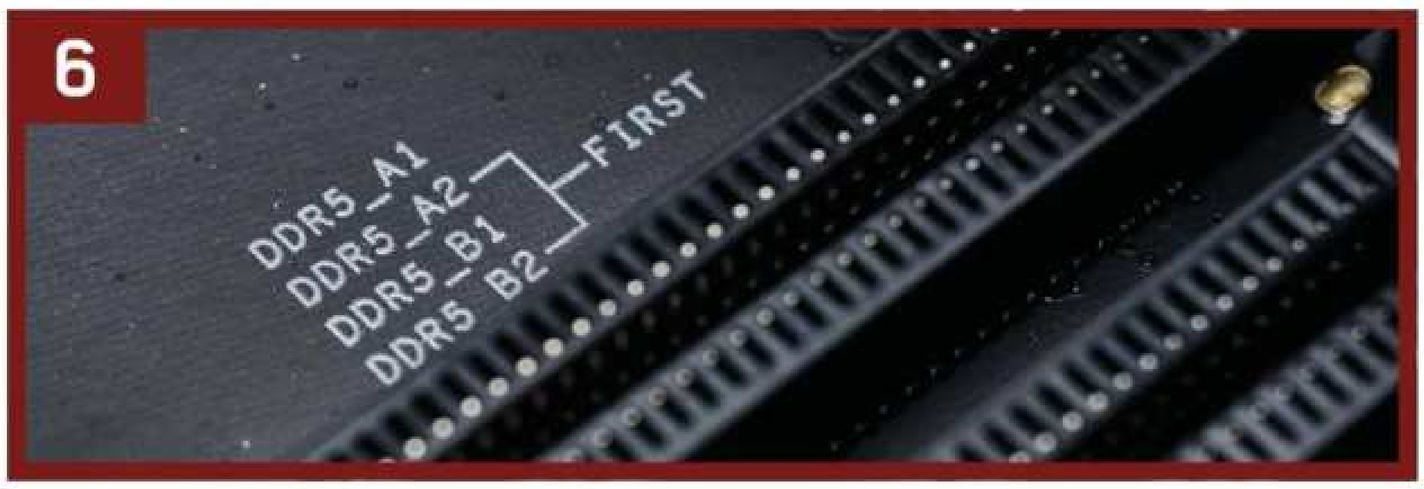

Fortunately, in the case of the H7 Flow 2024, almost everything is toolless, bar the front radiator bracket (it does have thumbscrews, but they're a little tight out of the factory). Pop off the panels, and place them in the chassis box for safe-keeping.

In our build, we foolishly requested the NZXT H7 Flow RGB variant instead of the stock standard one. That means it actually comes with a set of that Core Frame RGB fans as standard in the front of the chassis. That's fine, but given we're going with a high-contrast black and white build, it needs to come out. The front radiator bracket has two thumb screws at the top of it. Once the front panel is removed, you can loosen these off (either with sheer grip strength and force of will, or a good Philips screwdriver), then carefully lean the radiator bracket down and pull it out, the bottom of it sits in two slots **(Step 2)**. Lift it out, being mindful of the fan cable, and you can easily remove the frame by removing the four screws on either corner, then set it aside for later.

TEARING IT UP

With the front positively gutted, it's time to move to the rear of the case. NZXT's Flow has a 2.5-inch caddy that sits on the rear of the motherboard tray. Similar to the radiator bracket at the front of the chassis, it's held in place by a single thumb screw at the top, then notches at the bottom. Loosen it off, then slide it out **(Step 3)**. Word of advice: always save these brackets for later if this is a long-term rig—there's nothing worse than going to install a 2.5-inch drive at a later date, or needing to give it to a family member or friend, only for them to find that half the case is missing, and they can't buy a replacement.

Now our chassis is suitably prepped, it's time to move onto the motherboard itself. We've got a list of hardware that we can theoretically get installed right away. At this point, you should be able to install your M.2 SSDs, CPU, CPU backplate (potentially cooler, if you're running a fan setup), and your memory. Doing it outside of the case just makes life easier, particularly for the backplate.

Most modern case designs featured a pretty large backplate cutout on the motherboard tray, but trying to hold it from one side with one hand and secure it from the other is still finicky. In our build, we've gone down a little bit of a different route and opted to install the CPU and RAM after the motherboard's in the case directly. There's nothing wrong with doing it one way or the other, this was just more of a whim on the day kinda thing.

Anyway, NZXT's B650E for M.2 drive installation is super simple. The big block heatsink at the top has a single screw securing it. Undo that, take it out of its notch, and set it to one side. You're going to need two of the M.2 securing screws hidden away in the bottom of the motherboard tray as well for this. Then line up your primary OS drive, and place it in the top most slot, under that heatsink. Line up the cutout in the drive with the notch in the M.2 slot, and slide it carefully into place. Then push the drive down, and secure with the screw.

Next, remove the second magnetic plate cover below the PCIe x16 slot, and you'll spot another M.2 slot. This is where we'll be placing our WD SN5000 drive. It's the exact same installation as before. Once done, **(Step 4)** make sure you've removed the film from the thermal pads under both heatsinks, and secure them back into place. Fun nerd fact: if you take a look at that WD drive, you'll notice that it's got a single chip on top—that's all of its NAND Flash storage, and there's a ton of leftover space there on the PCB, meaning

STEP-BY-STEP GUIDE

there's potential for drives over 4TB if WD so desired. Neat, huh?

CALLING IN BACKUP

With the SSDs in, we moved on to the Kraken's backplate. NZXT has a pretty simple solution for AMD systems: remove the stock brackets one at a time, then install the stand-off plastic spacer and screw (in the correct orientation). It's all labeled in a bag marked 'AM5'. Once you've done one side, you can then do the other as well, and you're good to go—super easy and super simple **(Step 5)**.

One nice little detail on NZXT's B650E is this tiny DDR5 script etched into the PCB. No more hunting through little manuals trying to find out which slots you need to populate first for your DRAM; A2 and B2 are your friends **(Step 6)**.

This is where we deviated from the script slightly, opting to install the motherboard straight into the chassis. Grab the screws from the case accessories box (NZXT helpfully label them in tiny bags), then secure the motherboard in place. The N7 boards only have edge screws, so you don't have to worry about removing any plates to access hidden mounting locations underneath. Simply line it in position, then secure it down. Make sure the screws are nice and secure (not too tight, or you'll ruin the threads, and then end up with spinning stand-offs), and you'll be good to go. This also acts as a ground when you do this for your entire system.

With that, we also went ahead and installed our DRAM in the correct slots, as foretold by the NZXT PCB gods. Then, onto the CPU **(Step 7)** to install any AM5 (or LGA CPU for that matter), lift up the retention arm from the right side, then carefully raise the CPU bracket up. With that done, carefully lower the chip into position. There's a number of alignment tools at work here to ensure you always install the CPU correctly, from carefully placed notches and the branding on the chips themselves, to the tiny golden triangles on the corner. We almost always use the branding as a guide. Lower the chip into position so it rests carefully in the socket, give it a gentle wiggle with your finger to ensure it's secure, then replace the CPU bracket and its retaining arm, and you'll be all good to go. At this point, the cover should pop off if it doesn't pull it off by hand.

TOO EASY

With the motherboard now snug and prepped, we began work on setting up the cooling. Or our two fan grids? Arrays? Frames? Honestly, what do we call these things? Regardless, it's super easy. In the case of our radiator bracket from earlier, simply place your frame into position, secure from the other side onto the bracket, then mount the bracket back in the case. Pay close attention to where the cable is when you're installing it. Ideally, you want it going in that bottom right

cable hole. Then, do the same with the floor-mounted one, ensuring the fans are facing the correct orientation to draw cool air up and into the chassis itself. The fan grid should slot into position, then it can be secured from the other side (you will need to tip the case on its side to facilitate that). Again, make sure your cable is routing out to the back, and that you have six fans installed with eight screws, and some cables pre-run ready for attaching to the fan controllers **(Step 8)**.

Next up is the Kraken. Remove it from its packaging, immediately take the Intel bracket from the cooler, push it up towards the CPU display, rotate it slightly left or right, then pull it back down to remove. Then, take the AMD bracket, and do the opposite to secure it in position. Now, we can also pre-install the fans. Use the included screws that came with the radiator, and secure them in place on the underside of the radiator (where the tubes are facing down). From a performance perspective, push/pull makes very little difference, so aesthetics are the name of the game here. Similar to our other fan solutions, keep your cables in mind, ensure you think about the orientation of your radiator, how you're going to mount it in the case, then think about the best route possible for your fan cables to go out to the back of the case, then mount them.

Once done, lift the radiator (with the CPU block still loose) up into the chassis, and mount it to the ceiling with the included screws from the case accessory box. Then, remove the plastic cover from the CPU block itself, place it on the CPU (ensuring it's got ample thermal paste on it), secure the block into position, and you'll be all wrapped up **(Step 9)**.

A TRUE POWERHOUSE

With the vast bulk of the hardware now installed, it's time to move onto the power supply. Take it out of the box, and identify which cables you're going to need for your system. In our case, that's going to be the 24-pin, one PCIe power (daisy chained), two 8-pin EPS power (one of these will be split apart), and a SATA power for our fan controllers. Attach them to the power supply, then slide the PSU into position from the rear of the case, and once again secure with the included screws **(Step 10)**.

Here's another neat little addition, which we totally didn't spot until two thirds of the way through the build: the H7 Flow has an additional 2.5-inch caddy, hidden away at the front. Similar to the other bracketed items in the chassis, it's secured in place by a single thumb screw. Loosen that off and the HDD caddy will come loose. More cable management savings **(Step 11)**.

Now, it's time to move onto the proper cable management. We've gone ahead and installed both the 24-pin and the CPU EPS power. Running them up and through the various cable bars, velcro straps, cable channels, and grommets. NZXT really did revolutionize this with its initial line of H series cases, and the latest Flow is no different. Use the cable grommets to your advantage, and line everything up as best you can. The 24-pin will always be a bit of a nightmare to route. We've also pre-routed those AIO fan cables down, ready

STEP-BY-STEP GUIDE

to go into its own little fan controller. This in turn will then run into a motherboard header. We've also split off all of the front I/O cables ready for after this step. Cable separation is a weirdly helpful way of visualizing where all your cables need to go, and tidying up things. Break them up, and sling a bunch on top of your case off to the left, down under the PSU cover. Cable management is quite odd in a lot of ways—you don't necessarily need it to be flawless for a personal system, but doing the bare minimum like that makes it a lot less stressful (Step 12).

It's one of the things we always critique NZXT for with its PSUs, but the cables, for a $200+ PSU, are just a bit uninspiring, to be quite honest. There's no brading, flat cables, or black wires, sure, but this is something that be quiet! Corsair, EVGA, and co do far better (Step 13)—at least have an option for a pro cable kit, or something similar. PSU ports on the PSU side are often proprietary, so having the confidence to go out and buy a kit without worrying about it is a big deal.

CABLES MANAGED

With that wrapped up, we moved on to rear cable management. If you look at Step 14, you'll see how much capacity NZXT has implemented into the back of the H7 Flow. It's phenomenal. There are brackets you can wedge cables into, velcro straps everywhere, subtle channeling—the works. You'll notice that we've also installed the NZXT fan controller just north of the power supply there as well. That's specifically for the fans attached to the radiator. NZXT actually recommends that you place it top-left (there are literal markings etched into the case), and that's a fine place to put it, but just to keep the cabling a little tidier up that side, we opted to position it here instead. This is particularly handy, as we didn't have the 2.5-inch drive cage there, either.

We plugged the remaining fan frames directly into the NZXT RGB headers in the motherboard itself (Step 15). These are proprietary headers as far as we can tell, capable of handling the aRGB in NZXT fans directly, and run to a splitter, another of which needs to go into a fan header on the board. One of the minor limitations of that does mean that you can only run the fans within the frame at the same RPM, but honestly, over the years, the amount of times we've individually adjusted fans to operate slightly different to one another on one side of the case is zero.

Last but by no means least, we installed the GPU. As always, it's best to do this with the chassis laying on its back. We've also pre-routed the power cable through the cutout on the bottom. Then, after removing the two PCIe slot covers in the rear of the case, next to the PCIe slot we intend to use, we've slid the 7800 XT into position until it clicks into place, then secured it with the screws removed from those covers. With that done, it's time to plug the power in, and boom! All done (Step 16). At this point you're going to want to give your cables another go-around, making sure everything is tied down, and tidy. Then it's simply a case of replacing all those panels you removed earlier, and you're about ready to turn on this beast.

ROOM TO MANEUVER

THERE YOU HAVE IT, folks—one seriously slick mid-range(ish) PC built in one supremely svelte chassis. Generally, we always try to take around 16 images or steps for these builds. We actually struggled with our NZXT build, predominantly because everything went incredibly smoothly. There was very little that proved a challenge, or difficult to work around. There were no clearance issues—everything just worked. The cable management and the details and considerations made to tighten up the build procedure in the H7 Flow is just exceptional. There's little we can fault from a system build perspective at all.

It even gives you options. There's room here to install your AIO in the front if you so desire, and the liquid-cooling potential is immense, although you would have to think about how exactly to go about your reservoir and pump mounting solutions. But it's an impressively large amount of internal space on offer, for an incredibly minimal outlay. Similarly, the chassis itself really doesn't take up that much space on a desk either. It's your standard ATX Tower form factor, so to have such vast clearances for any and all hardware is, well, awesome.

If we were to approach the build again, going for a simple custom-cooling solution would definitely be our first thought. You've got capacity in here for three 360mm radiators in total (NZXT technically state that you can't mount a radiator in the floor, but we're not entirely convinced on that). Still, mount a 420mm in the front, 360mm in the roof, a long combo res on the front rad, and some simple yet durable black rubber soft tubing throughout, and this thing would sing—probably with something like a 7800X3D and an RTX 4090 to really get the most out of it.

But yeah, there's really little else to critique here. As far as the rest of the components go, it all lines up quite nicely. It would be great to see NZXT lean into a 'cable-less' AIO, and remove that ghastly cable running up and around the motherboard (or better yet, design some form of cable tidy on its boards directly), and of course, the power supply lacking better quality cables is a bit of a bummer, but that's really about it.

1 **Going for black and white parts might seem a little odd, but it adds a ton of contrast to the build, and really makes it pop, drawing your eyes to different elements in the system.**

2 **NZXT unfortunately doesn't have any sleeved cables available for its C series power supplies, so they're black and clean, but do look a little drab.**

3 **These triple fan frames are revolutionary. Honestly, if we could stick them in every build, we would. One cable out, four screws, and absolute control.**

4 **Speaking of fans, those bottom-mounted ones will help keep your GPU and overall system temps low.**

PERFORMANCE REVAMPED

HOW WE TEST

We've massively revamped our testing methodology for all of our builds moving forward. The aim is to provide a broad suite of benchmarks that consider how a system performs at both 1080p and 4K gaming, along with a range of benchmarks that push them to their absolute limits in traditional computational tasks, storage operations, and AI workflows. Additionally, we've expanded our testing to include power draw (utilizing a power meter and the same peripherals every time), alongside temperature monitoring throughout the benchmarking process.

This should give us greater insight into how these systems perform against one another, and better visibility, regardless of what the budget is, or how much more potent one GPU is over another.

THIS MONTH'S ZERO POINT

For our comparison build this month, we're utilizing a Frankenstein's Monster pieced together from July 2024's feature build. In that, it has the Core i9-14900K, the Asus ROG Maximus Z790 Dark Hero, 32GB of Lexar Thor OC DDR5 @ 6000 C36, a 2TB Crucial T700 M.2 PCIe 5.0 for our primary drive, and a 2TB Kingston Fury Renegade M.2 PCIe 4.0 for our secondary drive. The graphics card is Gigabyte's RTX 4080 Aero OC, and the power supply is the 1200W Corsair RMx Shift, all planted in the Geometric Future M4 chassis, and cooled by no less than six Phanteks M25-120 Fans, and a 360mm Glacier One 360D30 AIO. The total price for this build is $3,749 (54 percent higher), with a core price of $2,914 (111 percent higher).

RYZEN UP STRAIGHT TO THE TOP

Let's dive into our first new round of benchmarks. How does it do? Well, our AMD build is impressively efficient. Once upon a time, Team Red was almost branded as being super hot and super in efficient in how it operates. Today, that is absolutely not the case. A quick glance at the scores gives us a fairly good indicator on that front. Both Cinebench and Geekbench show that on a per thread basis, the Ryzen 7900X runs rings around the 14900K on pure efficiency alone. The scores are solid, too, with single-core performance only a smidge behind what Intel's best 14th-gen chip can muster.

Similarly, temperatures also held steady, topping out at 95.1 C, rather than Intel's flat 100 to achieve that. GPU temp was higher with the 7800 XT than the RTX 4080 Super, by around 5.2 C, but generally speaking system performance overall was fairly impressive. Storage also did pretty well. The SN5000, for a budget offering pumped out some decent numbers. Samsung's 990 Pro, was a little bit more underwhelming. $110 is a little on the steep side for a 1TB drive, but otherwise, it still did deliver 4K random read and writes close to that of the T700 in our Zero point.

On the gaming front, at 1080p, our AMD build nailed an impressive 117.86fps, although compared to the RTX 4080 and its 175.39fps, well, it's not really a fight. Similarly, at 4K, it managed 46.69fps as well on average. It's certainly possible to game at 4K with an RX 7800 XT, that much is clear, but some titles will struggle, particularly those with hardier ray-tracing elements, such as *Cyberpunk*. You can alleviate some of that with FSR, but it's still no where near as potent as what you'll find with Nvidia. DLSS 3.1 and AI frame generation is just outstanding when it comes to 4K. It close to tripled overall performance, bringing frame rates up from 27.4 to 77.3fps on average. The gap closes between the two cards marginally in *Total War*, with the 7800 XT landing 50.8 versus the 4080's 71.4, but it's still a sizeable lead for team green.

That's a theme that continues when we get into the 3D Mark performance, with *Speedway* in particular giving the RTX 4080 almost twice as much performance as its AMD counterpart. In *Steel Nomad*, the 4080 fared better, but not by much. Interestingly, in CPU Profile, the 14900K actually faired considerably better than the 7900X, landing a 15.3K score versus the 12.8K on the Ryzen, but again, it's close, and given the Ryzen 9 7900X comes in at $190 less than the 14900K, you have to question if it's worth it right now.

TEAM GREEN AI DOMINANCE

Then there's AI, and oh boy, what a shift. In these tests, we benchmark every component we can to provide a fair playing field. For the CPU, AMD clearly has an edge here, with the average inference time landing at just 520.37ms versus Intel's 677.14ms—the index scores support that as well. Jump to the GPU, and the Nvidia 4080 runs rings around the opposition, at least under Microsoft's ML API. The 4080 landed 9.645ms, compared to the 13.508 ms from AMD, with an index score almost double. Here's the thing: if we run that same test, but allow Nvidia to use its own TensorRT mode, it brings the time down to 6.539ms, landing a score of 2,690. Compared to the 50 and 56 the CPUs scored, the difference is astronomical.

For image generation, the results are even more in favor of Nvidia. Under ONNX, the RTX 4080 scored 2,585 with a 2.417 s/image generation time. Compared to the RX 7800 XT, it managed only 381, with an average time of 16.368 s/image. You can understand why the world's AI data centers are turning to Nvidia right now. If AI workloads are your bread and butter, and any of those workloads can take advantage of Tensor, Nvidia cards are the ones to consider.

Still, overall, given the price disparity, the AMD system puts in an impressive showing. It's considerably cheaper than its zero-point, yet in some cases comes close to matching Intel, certainly on the CPU-orientated tasks, and this is before the major micro-architecture changes launching in the middle of this month.

SYSTEM BENCHMARKS

For our system testing, we use a multitude of separate benchmarks. For our CPU tests, these include Cinebench 2024 and Geekbench 6.3.0. All results are indexes in this case. Multi-core per thread is achieved by dividing the total multi-core scores by the number of threads on the CPU. For storage testing, we're taking advantage of CrystalDiskMark 8.0.5. The best scores are in bold, and all results are reported in MB/s. Throughout our entire benchmarking process, we test each system with a power meter plug running, and the latest version of HWMonitor installed and up as well. This allows us to report on maximum power draw and temperatures seen over the entire benchmarking process.

CPU PERFORMANCE CINEBENCH 2024 & GEEKBENCH

	ZERO-POINT	
Cinebench 2024 Multi Core	1,902	1,592 (-16.3%)
Cinebench 2024 Multi Core Per Thread	59	66 (11.9%)
Cinebench 2024 Single Core	124	118 (-4.8%)
Geekbench 6 Multi Core	20,044	18,155 (-9.4%)
Geekbench 6 Multi Core Per Thread	626	756 (20.8%)
Geekbench 6 Single Core	2,934	2,919 (-0.5%)

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

STORAGE PERFORMANCE CRYSTALDISKMARK 8.0.5

	ZERO-POINT	
C: Sequential Read / Write Q8	11,974 / 11,566	7,435 / 6,871
C: Random 4K Read / Write Q1	85.75 / 304.90	77.05 / 289.64
D: Sequential Read / Write Q8	7,039 / 6,721	5,531 / 5,029
D: Random 4K Read / Write - Q1	80.63 / 328.98	75.70 / 248.03

POWER DRAW & TEMPERATURE HWMONITOR & DEDICATED POWER METER

	ZERO-POINT	
Max CPU Temperature	100.0	95.1 (4.9%)
Max GPU Temperature	73.2	78.4 (-7.1%)
Max System Power Draw	649.9	509.3 (21.6%)

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

GAMING BENCHMARKS

Game tests are performed on 'Ray Tracing Ultra (*Cyberpunk*)', 'Ultra (*Far Cry 6*)' / 'Maximum (Final Fantasy XIV Dawntrail)' and 'Ultra (*Total War: Warhammer 3*)' graphics presets respectively. Three runs are performed per title, per resolution. *Cyberpunk* is tested with ray tracing enabled, regardless. Under its second set of conditions, DLSS/FSR/XeSS is enabled in *Cyberpunk*, dependent on GPU manufacturer, with upscaling set to 'Quality'. Nvidia cards also have 'AI Frame Generation' enabled. *FF XIV Dawntrail* is utilizing FSR across the board regardless of GPU, with 'LOD on Distant Objects' enabled and dynamic resolution disabled. *Far Cry 6* is tested with DXR disabled and the HD Texture pack downloaded and enabled.

1080P PERFORMANCE

	ZERO-POINT	
Cyberpunk 2077 Stock (Average fps)	96.1	44.2
Cyberpunk 2077 Stock (Minimum fps)	82.9	38.1
Cyberpunk 2077 AI Upscaling (Average fps)	207.2	73.0
Cyberpunk 2077 AI Upscaling (Minimum fps)	164.3	62.8
Far Cry 6 (Average fps)	169.7	149.7
Far Cry 6 (Minimum fps)	131.0	104.7
Final Fantasy XIV Dawntrail (Average fps)	231.0	174.4
Final Fantasy XIV Dawntrail (Minimum fps)	94.2	77.7
Total War: Warhammer 3 (Average fps)	173.0	148.0
Total War: Warhammer 3 (Minimum fps)	123.3	113.3

1080P PERFORMANCE

	ZERO-POINT	
1080p Titles Aggregate (Average fps)	175.4	117.9 (-32.8%)
1080p Titles Aggregate (Minimum fps)	119.2	79.3 (-33.5%)

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

4K PERFORMANCE

	ZERO-POINT	
4K Titles Aggregate (Average fps)	76.6	46.7 (-39.0%)
4K Titles Aggregate (Minimum fps)	52.0	37.0 (-28.8%)

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

GAMING BENCHMARKS *CONTINUED*

4K PERFORMANCE

	ZERO-POINT	
Cyberpunk 2077 Stock (Average fps)	27.4	13.1
Cyberpunk 2077 Stock (Minimum fps)	18.5	11.4
Cyberpunk 2077 AI Upscaling (Average fps)	77.3	25.6
Cyberpunk 2077 AI Upscaling (Minimum fps)	40.4	22.3
Far Cry 6 (Average fps)	109.7	83.0
Far Cry 6 (Minimum fps)	88.7	76.0
Final Fantasy XIV Dawntrail (Average fps)	97.0	61.0
Final Fantasy XIV Dawntrail (Minimum fps)	60.8	39.3
Total War: Warhammer 3 (Average fps)	71.4	50.8
Total War: Warhammer 3 (Minimum fps)	51.7	36.0

SYNTHETIC GAMING BENCHMARKS

For our synthetic graphics performance tests, we're utilizing 3DMark's latest Speed Way, Steel Nomad, Night Raid, and CPU Profile tests. Speed Way represents a DX12 title at 1440p, with ray tracing elements. Steel Nomad is DX12 4K without ray tracing, and Night Raid (DX12) is designed with integrated GPUs in mind at 1080p. CPU Profile acts as an additional CPU benchmark (reporting on Max Thread output only), but with specific gaming workloads prioritized. Each synthetic test is run a total of three times, and Indexes are reported on. In the case of Night Raid, only the overall score is listed. However, we do keep logs of Graphics and CPU as well to identify any erroneous elements or bottlenecks.

3DMARK PERFORMANCE

	ZERO-POINT	
Speed Way (DX12 + RTX @ 1440p)	7,241	3,760 (-48.1%)
Steel Nomad (DX12 @ 4K)	6,491	4,024 (-38.0%)
Night Raid (DX12 @ 1080p)	94,118	83,131 (-11.7%)
CPU Profile	15,269	12,776 (-16.3%)

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

AI INFERENCE BENCHMARKS

For our AI testing, we're utilizing UL Procyon's Professional Benchmark Suite. The ones we're focused on in particular are its AI Computer Vision benchmark, which gives us a gauge on how our systems operate when dealing with inference engine tasks, and AI Image Generation benchmarks. For Computer Vision, we report on the number of inferences performed, the average inference time, and the Index score across all tests. The Index score is more valuable, as it's the time taken between Inferences that indicates performance. We run all tests on AI Vision utilizing Float16 precision. We test both the CPU and GPU under Microsoft's ML API as standard, and if viable, include figures for testing the GPU under Nvidia's TensorRT SDK, Intel's OpenVINO toolkit, or Qualcomm's SNPE run time (specifically for ARM devices), depending on the comparison zero-point.

For Image Generation, we test under Stable Diffusion 1.5 (Float16 precision) using all engines available. However, we'll only ever include scores from the same engine. We report on the Index score, the time taken, and the image generation speed. All results are averaged across all tests. The best scores, where appropriate, are in bold.

AI COMPUTER VISION PERFORMANCE

	ZERO-POINT	
Microsoft ML API - CPU (Index)	50	56 (12.0%)
Microsoft ML API - GPU (Index)	1,945	1,084 (-44.3%)

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

AI COMPUTER VISION PERFORMANCE

	ZERO-POINT	
Microsoft ML API CPU (Total Inferences)	31,474	11,278
Microsoft ML API CPU (Average Inference Time)	677.14 ms	**520.37 ms**
Microsoft ML API GPU (Total Inferences)	685,207	136,607
Microsoft ML API GPU (Average Inference Time)	**9.645 ms**	13.508 ms

AI IMAGE GENERATION PERFORMANCE

	ZERO-POINT	
Stable Diffusion 1.5 ONNX (Index)	2,585	381 (-85.3%)

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

AI IMAGE GENERATION PERFORMANCE

	ZERO-POINT	
Stable Diffusion 1.5 ONNX (Time Taken)	**38.676 s**	261.885 s
Stable Diffusion 1.5 ONNX (Image Generation Speed)	**2.417 s/images**	16.368 s/image

OVERCLOCKING IN 2024 HAS THE END ARRIVED FOR HIGH VOLTAGE?

Zak Storey takes a look at whether overclocking is really worth it anymore

THERE'S SOMETHING inherently human about pushing technology to its limits. Companies do this all the time: building, developing, iterating on designs, new hardware, trying to eek out every last fleck of performance to beat the competition. The same can be said for amateur enthusiasts. There are very few people that don't try to optimize how they work, their PC, or the hardware inside it.

Overclocking, by its very nature, speaks to that. In essence, it's a love language for your PC. In the same way that you can get attached to a car, tool, or fishing rod, you can also become attached to a processor. It's a bit of a bizarre quirk of humanity, this attachment to objects, as if they have personality; a soul.

The thing is, some processors actually do, to a lesser extent. The silicon lottery gifts them with differences. Some are better than others; some worse; some perform in a particular manner if you do one little tweak, and some score higher in one test than an identical CPU at the same spec. Apply enough voltage, adjust the multipliers, increase those clock speeds, and you could draw out more performance for free—at least that's how it used to be. So what's changed? Has the golden era come to an end? In a world of high competition and with the likes of Arm breathing down Intel and AMD's necks, has that at-home pastime retired to the humble shores of PC history? Let's take a look.

OVERCLOCKING'S HUMBLE ORIGINS

It might come as a surprise, but overclocking has been around for an incredibly long time. In fact, theoretically speaking, you could very much argue that the act of tweaking clock speeds actually originates with the OG manufacturers of those CPUs as they dive into trying to find exactly how far they can push those clock-speeds while keeping the chips themselves stable.

Overclocking in the manner that we are perhaps more accustomed first landed with us in the early '90s when Asus and a select few other manufacturers debuted the first motherboards that gave you the opportunity to tweak those speeds. Asus in particular had its i386 motherboards that allowed you to adjust the multiplier, CPU speed, and voltage simply by tweaking a few jumpers on the motherboard itself. This was very loosely overclocking; you had preset values and voltage settings that you could apply. The silicon lottery was still a thing, even back then, so depending on your luck, it would actively affect what performance you could get out of the chip.

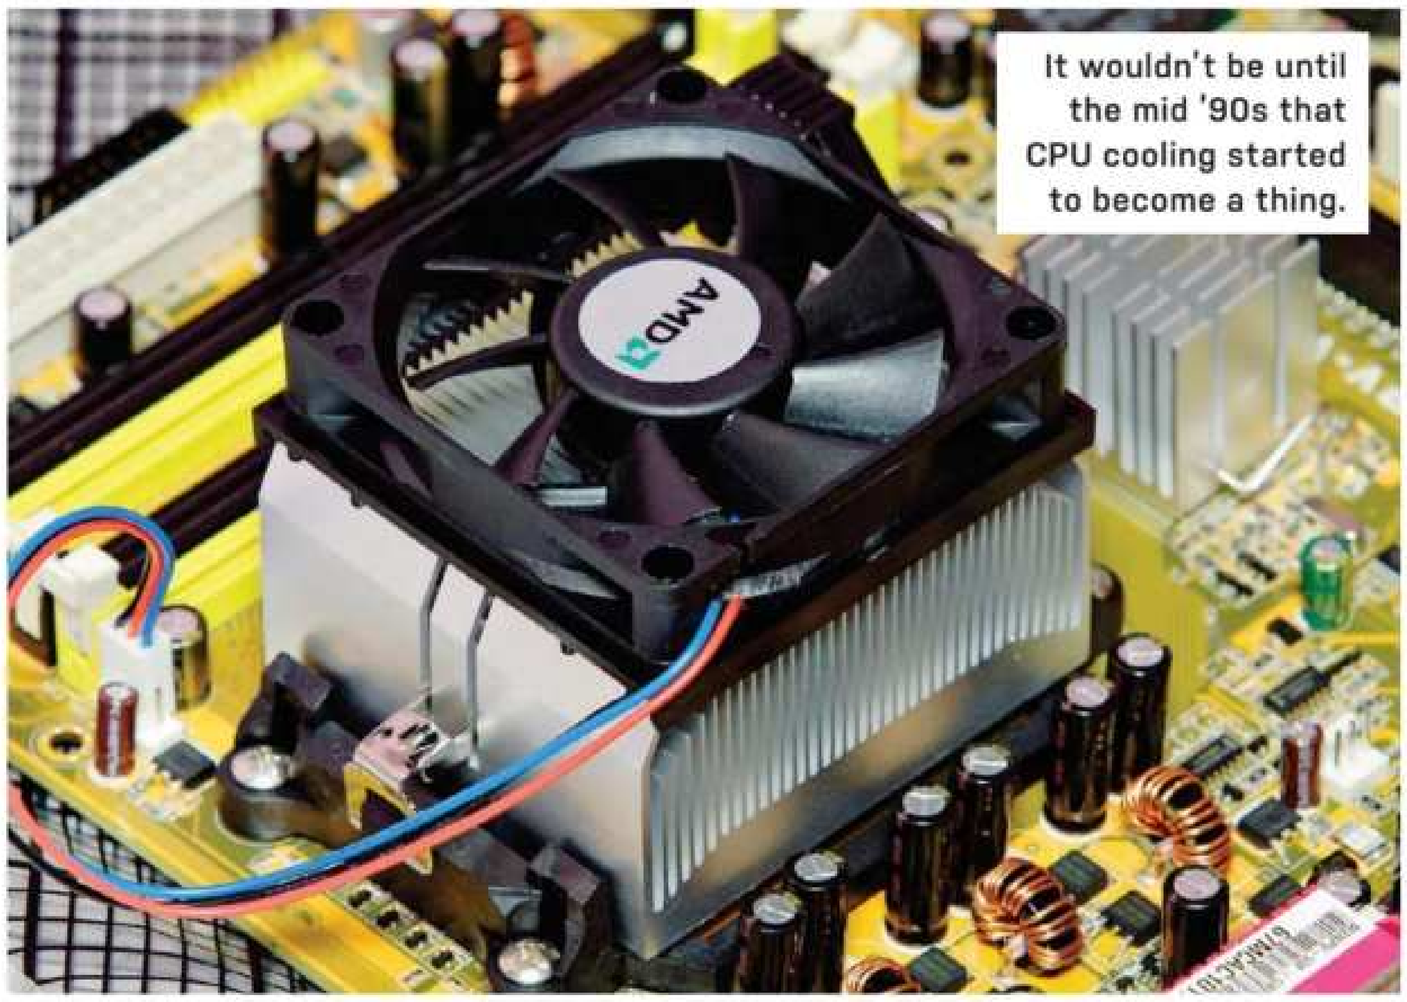

It wouldn't be until the mid '90s that CPU cooling started to become a thing.

© WIKIPEDIA/ FIR0002/FLAGSTAFFOTOS

THE ADVENT OF COOLING

All of this was, of course, done without direct cooling on the CPUs themselves. Specific CPU air towers and other such solutions wouldn't arrive until slightly later on, designed to alleviate that overclocking heat. Similarly, this is also about the same time that the first extreme overclockers arrived on the scene, complete with dry ice and LN2.

The first unlocked chips landed in the early 2000s with Intel's Pentium series of processors. Finally, this gave power users the opportunity to tweak with their chips in a similar manner to what we'd see later on in the industry. Likewise, overclocking tournaments and competitions began to blow up as more and more manufacturers pushed cash and resources into the sport in an aim to take the accolade as the motherboard manufacturer that managed the world's highest frequency.

THE AMD PENCIL TRICK

Back in the early 2000s, it was possible to overclock AMD Thunderbird and Duron processors by using a pencil to fill in the contacts on the L1 Bridge. These tiny contacts limited the CPUs to only operate at a certain frequency straight out of the factory. With them penciled in, you could easily achieve 650-700 MHz on the Durons, and even as high as 900MHz on the Thunderbird chips.

AMD did try to stop end-users from doing this on its next generation processors by applying a layer of varnish over those points. However, enthusiasts managed to bypass this by using a silver lacquer paint instead, again overclocking their processors by a considerable margin. As long as the cooling was sufficient, you could then adjust the multiplier in the BIOS to get those speeds.

THE CORE SERIES

In the late 2000s, moving into the 2010s, Intel's Core series finally debuted, and with it overclocking was now mainstream. The humble BIOS had been redeveloped and revamped to allow for mouse control, and all of a sudden, it was far easier to adjust and alter multipliers and voltages quickly and without too much hassle. Simply jump into the BIOS, adjust the multiplier, change the voltage, benchmark in Cinebench or Prime95, and you were good to go. Failed attempt? Tweak the settings, and try again. With the now ubiquitous use of the internet, and a whole host of resources breaking out into the mainstream overclocking was easily achievable by the vast majority of system tinkerers.

THE MODERN ERA AND AUTO-OVERCLOCKING

Over the years, overclocking remained relatively the same. While processors often came with their own advanced forms of Turbo, these were applied on a wide spectrum, and were aimed more at ensuring processor stability, rather than perfectly tuned to each individual chip—all vulnerable to the silicon lottery's whims.

Interestingly, even Nvidia's GPU Boost, when it first debuted with its Kepler 600 series graphics cards, likewise followed a similar routine, where the card would ramp up its clock speeds, dependent on whether or not it was close to utilize its full TDP, and within temperature limits. In fact, this was an easy way to auto-overclock your graphics cards back in the day (and even today, to some extent) simply by using something like MSI's Afterburner to slide the 'power target' all the way up to its maximum, in some cases by as much as 125 percent. That would give the graphics card the opportunity to boost its clock speeds even higher.

Generation on generation, however, these technologies became increasingly advanced, working off both power draw and temperature on CPU and GPU respectively. With the likes of AMD's Precision Boost, Extended Frequency Response, and Precision Boost Overdrive, working hard to ensure that clock speeds remained higher for longer periods of time during heavy workloads if the cooling, power delivery, and stability were within acceptable parameters.

MARKET CONDITIONS

Here's the thing, though: over the last few years and generations of processors, we've seen both AMD and Intel continually push their chips to the absolute limits. Binning processes have

© WIKIMEDIA/ DIRK OPPELT

Looks like that power target slider does little in regard to real world performance.

become more complex in an effort to weed out inferior chips, and auto-overclocking tech has become more advanced, with even motherboard manufacturers getting in on the action to try and give themselves a competitive edge over each other and ensure you buy their products.

In a lot of ways, that is a good thing. However, as Moore's law diminishes, and we start to reach the physical limits of what's possible when it comes to chip performance entirely, overclocking is very much starting to feel like a thing of the past. LN2 and extreme overclocking still exists, as electrons behave very differently at those kinds of temperatures. However, for modern hardware tinkerers, and with today's CPUs being so aggressively amped up, the question is, what's the point? Particularly when we're seeing stability issues and potential hardware damage occur on modern-day chips, and that's just because of incorrect application of voltages provided by motherboard manufacturers, and Intel themselves.

SO, WHAT'S A MODERN OVERCLOCKER TO DO?

Well, here's the thing: the majority of overclocking today very much relies on just enabling the settings. PBO (Precision Boost Overdrive), for instance, is a toggleable feature that you can configure directly in the BIOS itself for all Ryzen 9000X chips. How much extra performance does that deliver, though? Well, from the limited testing we've done, on average you're looking at a performance increase of three to four percent, dependent on application, certainly for multi-core performance, and around one percent for single core.

Similarly, if you're thinking that Intel might have an edge here, you'll be sadly mistaken. In the case of the 14700K, even with our best attempt at an overclock, performance in Cinebench 2024 increased by a mindblowing 0.8 percent on the multi-core test, and even less than that on single-core.

There are auto-overclocking solutions as well, typically baked into both Intel and AMD's own overclocking software. Intel has its Extreme Tuning Utility (or XTU) and AMD has its Ryzen Master, both of which features 'automatic overclock' buttons. In both scenarios, Intel and AMD stress test the CPU in the machine, before applying and tweaking the voltages, core ratios, and more, to find a stable overclock specific to your chip.

With Intel's 14th generation, the auto-overclocking feature is actually baked around AI machine learning models, as it looks

MAXIMUM PC FACTS WITH ZAK


The build that started it all.

It's a little known fact, but the reason Zak landed his first job at *Maximum PC* as a staff writer was partly because of his ability to overclock. In his portfolio, not only did he include a CV, cover letter, writing exercise and extensive photography of the build's he'd produced, but also screengrabs of an Intel Core i5-2500K overclocked to a solid 4.6 GHz.

A fairly reasonable overclock for the 2500K, what made that even more impressive, however, was the fact that it was managed on an ASRock Z77E-ITX motherboard, inside a Bitfenix Prodigy chassis, with nothing but a 120mm Corsair H60v2 CPU and a couple of fans, and still managed to keep itself cool at just 65 C.

In fact. you can find Zak's first ever feature, an 'overclocking masterclass' in the September 2015 issue, off the back of that one little overclock. Who said tinkering doesn't get you places?

Qualcomm's Snapdragon processor is revolutionary, and marks the probable end of x86 as we know it.

at how your chip reacts under load, and then applies custom settings to it based off the results it initially gets from those stress tests. That's impressive, as it can manage all of that in less than a minute. Again, this still very much voids the warranty on your processor if you decide to go down the XTU route.

Similarly, GPU overclocking is now pretty much dead. Although you can adjust memory overclocks and frequencies on the latest cards, it's nowhere near as easy to do as on the previous generations. Almost all modern cards feature an ECC-type element to the memory stacks, ensuring that they recalculate any corrupt memory states during operation. This means that to actually overclock your memory using something like MSI Afterburner, you have to pay attention to your overall performance. Overclock just your memory, then wait until you see a marked decline in performance the higher you go. Similarly, most modern GPUs sit within their power limit parameters. In fact, Nvidia no longer advertises GPU Boost in its cards, as it's an inherent part of a modern GPU. Still, you can adjust the power target in GPU control software, although again, the difference in overall performance, even with slightly higher clock speeds, is relatively negligible in real-world scenarios.

As an example, with the power target ramped up on a Gigabyte RTX 4080 Aero OC, in *Cyberpunk* at 4K on the Ray Tracing Ultra Preset, with DLSS set to Quality and Frame Generation on, we got 76.37fps on average and 40.43 min. Max clock speed topped out at 2895MHz, and GPU temp hit 70 C in the short time it was running. At stock, the stats read pretty much exactly the same (the temperature actually rose by 1 C, likely ambient), plus the average frame rate actually 'increased' to 76.48 (0.14 percent) and minimum frame rate increased to 41.05 fps (1.53 percent).

WHY CLOCK SPEED REALLY DOESN'T MATTER.

The reality is that we've reached a point in general processor design where clock speed isn't a good indicator of overall performance. We've seen this with pretty much every single product stack. AMD has pivoted to 3D cache to improve its CPU performance, without necessarily increasing the number of cores on any of its mainstream chips. Intel has opted to go for a dynamic core system with both performance and efficient systems, and of course, we can't not mention both AMD and Nvidia, who are now banking hard on technologies like AI supersampling and ray tracing to really push the limits of what's possible.

We can't forget Arm processors either, utilizing an entirely different processing method via RISC to push the boundaries of what's possible when it comes to efficient processing.

Comparing clock speeds between Intel, AMD, and Qualcomm processors, given how dramatically different their core architectures are, and how they operate at a base level, is practically fruitless, and in no way an indicator of which is better. The sad truth is that overclocking may have reached its end. What little performance there is to be had just isn't the worth the extra instability, temperatures, and power draw. True performance boosts are seemingly coming from dedicated co-processors, and new avenues on how to compute specific tasks more efficiently.

Even memory suffers from the same problem. Although clock speed is a good indicator of memory speed, in real-world scenarios, it's actually the latency and connection between the CAS latency and memory speed that dictates performance. Memory manufacturers often struggle when a new generation of memory is introduced, as the increased clock speeds and capacity often come along with increased CAS latencies, reducing real-world performance by an impressive margin.

The sad reality is that no matter which way you cut it, the death of overclocking seems nigh. Long gone are the days where you could take a consumer processor, draw a few graphite lines on it, and boost its performance by 25-30 percent. Even the best overclocks we have today pale in comparison to that, and you'd be lucky to even manage a three to four percent performance boost. It's like a piece of history has passed us by, lost to obscurity with little fanfare. On the other hand, with performance now arising from more dynamic inventions and engineering, the future's never been so exciting as it is now.

UNDERVOLTING

Don't judge those tubing runs, we were new!

There's another side to overclocking, and that is undervolting. This pursuit is a little different compared to just ramping up clock speed, volts and temperatures in that it generally involves reducing voltage from the chip itself, while trying to retain its performance and clock speeds as advertised. Now, why the heck would you do that? Admittedly, it is an incredibly niche scenario, but it's often used in noise-sensitive systems that require absolute silence, or as close to as you can get.

We did this with a liquid-cooled PC inside the NZXT Manta back in July 2016. The system featured an AMD Radeon Fury X alongside an Intel Core i7-6700K, one that was pre-binned by a third party, and capable of hitting 4.9GHz. Because the thermal capacity of the chassis was quite limited, instead of leaving it running at stock or overclocking, and having sky-high temperatures, we opted to go the other route, and drop the volts while retaining those clock speeds. This meant we could keep those IPPC Noctua Industrial fans running at lower RPMs, reducing overall noise under load. To do this, we reduced the VCore voltage in a similar manner to how you'd overclock a CPU, but in the other direction. We did this until the chip became unstable, or the clock speed began to drop.

By the end, we had an Intel Core i7-6700K, blasting through benchmarks at the same speed as its stock settings, but with a maximum temperature under load of just 47 C. How times have changed.

magazinesdirect.com

Over 100 brands to choose from

GIFT IDEA! Gifts that keep on giving

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 October 31 2024.

BUILD AN IT SUPPORT HUB

Discover how to use RustDesk to provide remote assistance and control your own devices remotely with *Nick Peers*

AS A MAXIMUM PC READER, your expert knowledge of PCs no doubt makes you a popular target for technical support requests from friends and family. "It's why I sent you to university," this author's mother used to say, forgetting he studied history and that he fell into technology journalism by happy accident.

Trying to give tech support over the phone or by email is a nightmare if you're constantly having to explain the basics or trying to imagine what's on the screen of the ailing machine. This is where remote control software comes into its own—Microsoft's Quick Assist tool is one option, but you're tied to offering PC-to-PC support and it can't handle administrative requests, which is useless when trying to access certain settings or install fix-it tools.

In the past, the go-to tool for remote assistance—as well as giving you access to your own home devices when out and about—was TeamViewer. It's simple to use and works across all major platforms, including macOS, Linux, and Android, as well as Windows. However, its free tier has been locked down to the point of being almost unusable as you find yourself unfairly accused of using it for professional purposes.

The good news is that we've found an alternate option offering all the advantages of TeamViewer with none of the charges, nags, or restrictions. Read on to discover how to set up and use the brilliant RustDesk.

Use the chat window to communicate with the person you're helping.

HOW RUSTDESK WORKS

We've chosen RustDesk (https://rustdesk.com) because it is free, open source, and simplifies switching thanks to its familiar-looking user interface. It also has a built-in file transfer tool that makes it easy to send and receive files between devices.

RustDesk's client package can be used to remotely access and control Windows, Mac, and Linux computers (including ChromeBooks with Linux Developer Mode enabled), as well as Android devices. There's also a client app for iOS that enables you to control other devices from your iPhone or iPad. The client is all you need to install on both your PC and any remote device.

By default, RustDesk connects through a free public server, just as you would with TeamViewer. All you need to know is the device's ID and—if the remote machine is unattended—a one-time password that's randomly generated each time the client is launched.

One of the issues with this approach is that performance is noticeably slower than TeamViewer, due to RustDesk's less developed infrastructure. However, you have two additional connection options available to you: directly through IP address (for local connections), or through your own RustDesk server, which helps improve performance. We'll cover all three options in this feature.

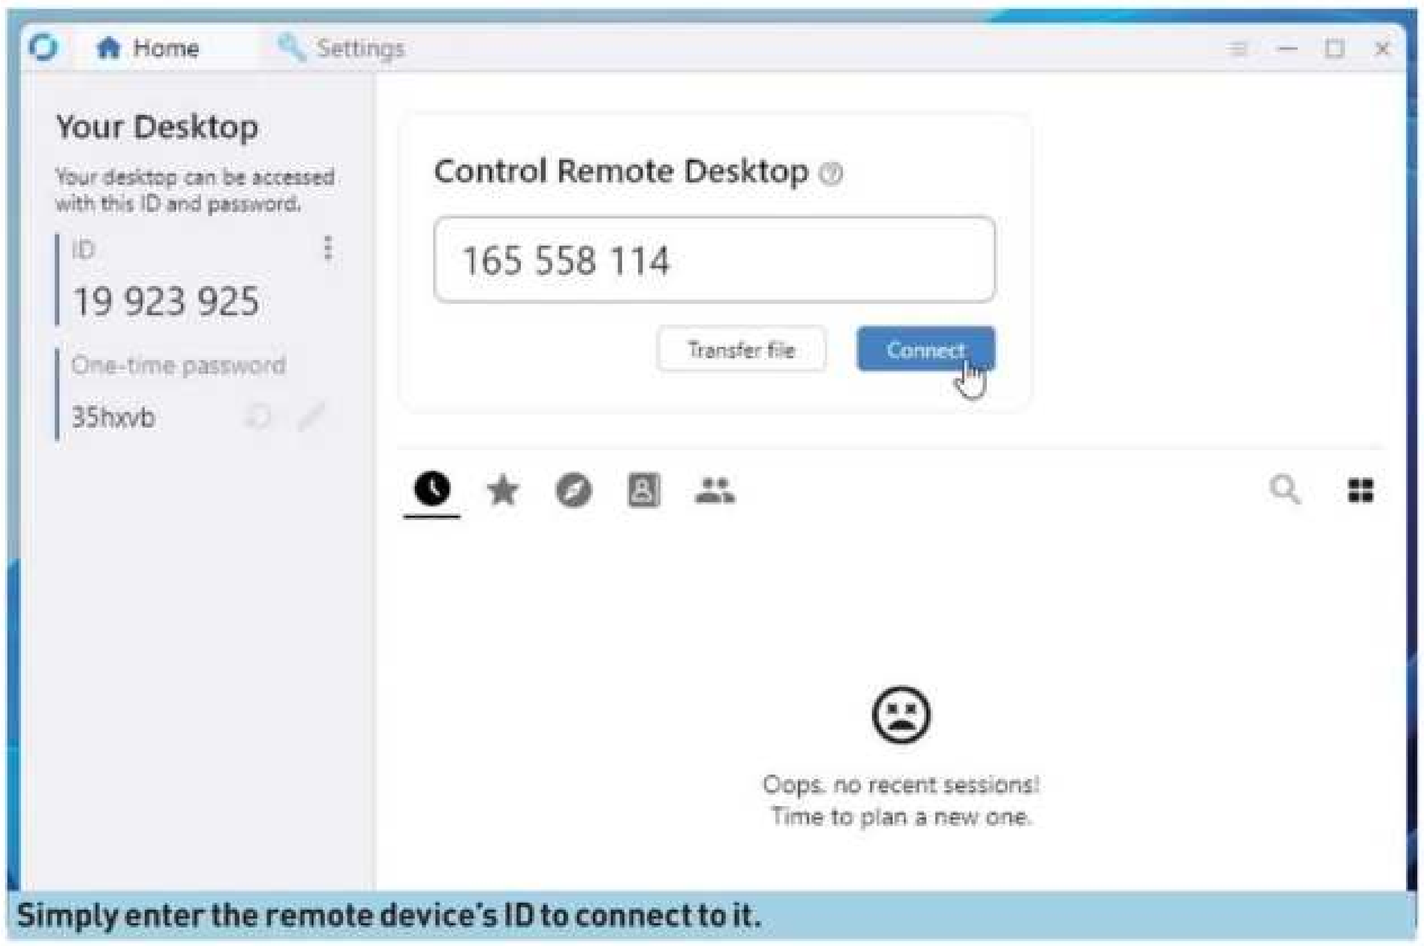

Simply enter the remote device's ID to connect to it.

FIRST STEPS

Whatever approach you take, the RustDesk client must be installed on both devices to handle both sides of the connection. Head over to www.rustdesk.com/download where you'll be rerouted to the RustDesk client GitHub page, complete with a convenient list of links spanning all supported devices: Windows, Mac, Linux, Android, and iOS. The iOS version is the only one that can only be used to access other devices.

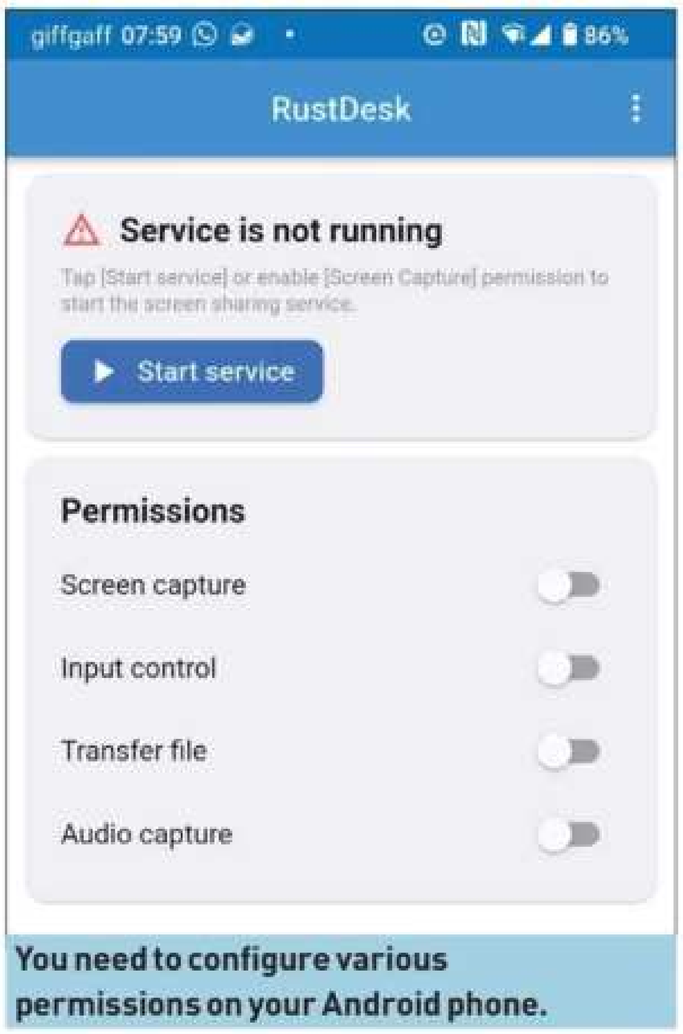

You need to configure various permissions on your Android phone.

You'll find direct links for all platforms here—including APK files for Android devices, because RustDesk isn't available in any app store. Once installed, the client automatically opens. Some platforms (such as Mac) may then require you to configure additional permissions to give RustDesk the access it needs—just follow the prompts.

Android users also need to jump through several hoops to configure the app to allow the device to be controlled remotely: First tap 'Start Service', then read the warning message about scams. Wait for the 'I Agree' box to turn blue, then check 'Don't show again' before tapping it. You're then prompted to configure other permissions, such as 'Display over other apps' and 'Files access' for transferring files. You then appear to run into a brick wall should you want to control your Android device remotely, because tapping 'Input control' shows a grayed-out 'RustDesk Input' option.

Don't worry—this appears because you installed RustDesk from an APK file, and the fix is mercifully simple: Open Settings and navigate to 'Apps > RustDesk', then tap the 'More' button (the vertical ellipsis) at the top and select 'Allow restricted settings'. Now return to Accessibility, where you can now give RustDesk Input full control of your device.

MAKE THE CONNECTION

The main RustDesk client screen contains everything you need to both receive and make connections. If you're connecting through RustDesk's own public servers, you need to know the ID of the device to which you're connecting. If you're connecting to your own device, you also need to know the password; if you're connecting to a friend or family member who's sitting at their device, they can grant you access when you attempt to connect.

To make your first connection, type the remote ID into the 'Control Remote Desktop' box at the top of the screen and hit 'Connect'. This triggers an access request pop-up on the remote device—be patient, this can take some time over RustDesk's public servers. Either enter the password and click 'OK' or wait for the person at the other end to click 'Accept'.

This approach ensures control remains in the hands of the device owner. If it's your own device, you'd simply enter the password to gain control of it; if it's a friend or family member looking for help, they can either share the password with you beforehand, or sit at their device and click 'Accept' when the prompt appears on screen. See the 'Set access permissions' box on the right for more details about how they retain full control over their device, including being able to terminate the connection at any time.

NAVIGATE THE REMOTE DEVICE

When your connection is established to the remote device, you see the remote desktop appear in a pop-up window on screen. If you have control over the keyboard and mouse of the remote device, you can now use that device as though you were physically sitting next to it.

By default, RustDesk attempts to render the screen at the same resolution as the remote device—in many cases,

You can access Macs as well as PCs (Windows and Linux).

you'll find it doesn't fit within the confines of the RustDesk client window, so only part of it is shown. Consequently, as you roll your mouse over the window, you see the desktop move with it.

While this works, it's not very practical—particularly on target machines with a higher resolution screen than your own (say 4K on a regular 1080p HD display). The good news is that there are plenty of options for changing the display to suit.

Look for the small tab at the top of the window displaying three buttons—reposition the tab by clicking and dragging on the gray left-hand button. The middle button enables you to switch RustDesk to full-screen view—perfect for machines where both host and remote displays match (you'll see the tab remains in view at the top, with buttons for exiting back to normal view or minimizing the window to the taskbar).

But what about 'fitting' a larger-resolution remote display to your RustDesk window? To do this, click the right-hand ˅ button on the tab to reveal seven icons. Click the display icon (third from left) to reveal a pop-up menu. You'll see 'Scale original' is selected by default, so choose 'Scale adaptive' to have the remote display fit to the window. Alternately, expand the Resolution submenu to reveal options for setting the remote device's resolution to a more comfortable setting if you prefer. For more display settings, check out the 'More display options' box over the page.

INPUT DEVICE SETTINGS

Next to the display button is a keyboard button—click this to reveal various options for configuring your input device. The key option is 'View mode'—when checked, it basically turns the session into a simple screen-sharing one by disabling access to the remote device's mouse and keyboard while graying out all other options. When you're in control, you'll see options referring to various modes: Legacy, Map (the default), and Translate. Ignore Legacy—this would only be applicable if you were attempting to connect to a device running an older version (1.1.9 or earlier) of the RustDesk client.

Instead, consider switching to Translate if you want to do office work on the remote machine, because this mode effectively applies your local keyboard layout to the remote machine, ensuring that what you type on your local keyboard is input at the other end, too (for example, ensuring you don't fall foul of a remote Mac's transposition of the " and @ characters). Note, Translate mode isn't recommended for gaming, where scan codes may be used to process keypresses.

Other input options are largely self-explanatory: 'Swap control-command key' helps circumvent the Mac's insistence

SET ACCESS PERMISSIONS

When rendering technical assistance to a friend or family member, ask them to carefully study the access request pop-up before clicking 'Accept'. By default, the seven permissions buttons are all colored blue to indicate they're all enabled. This effectively grants you full control over their device, so they may wish to review these buttons and then click any they're uncomfortable with to turn the button gray and so disable it.

The first row of buttons gives you control over their mouse and keyboard, plus access to their clipboard and audio output (so you can hear what's going on). There's also support for file copy/paste, which would enable you to transfer files to their device without having to go through the file transfer tool.

The second row enables you to remotely restart their computer if necessary, set up a recording session that would record the session as a video file, plus give you the option of blocking user input from their end for whatever reason (typically because you don't want them interfering as you try to fix their problem).

Ultimately, the most important button for remote control is the first one: Disable the mouse and keyboard, and you can't interact directly with the remote device, effectively turning it into a simple screen-sharing session.

The remote device owner can change permissions and end the connection at any time.

Note, once your friend clicks 'Accept', the pop-up is minimized to the taskbar—your friend can bring it back to change permissions at any time, plus they can terminate the connection by clicking 'Disconnect'. Be aware that while you can move and hide this window, you can't change any of the settings from your end. If you wish to end the connection, click the X in the RustDesk window to close it and return to the main RustDesk connection screen.

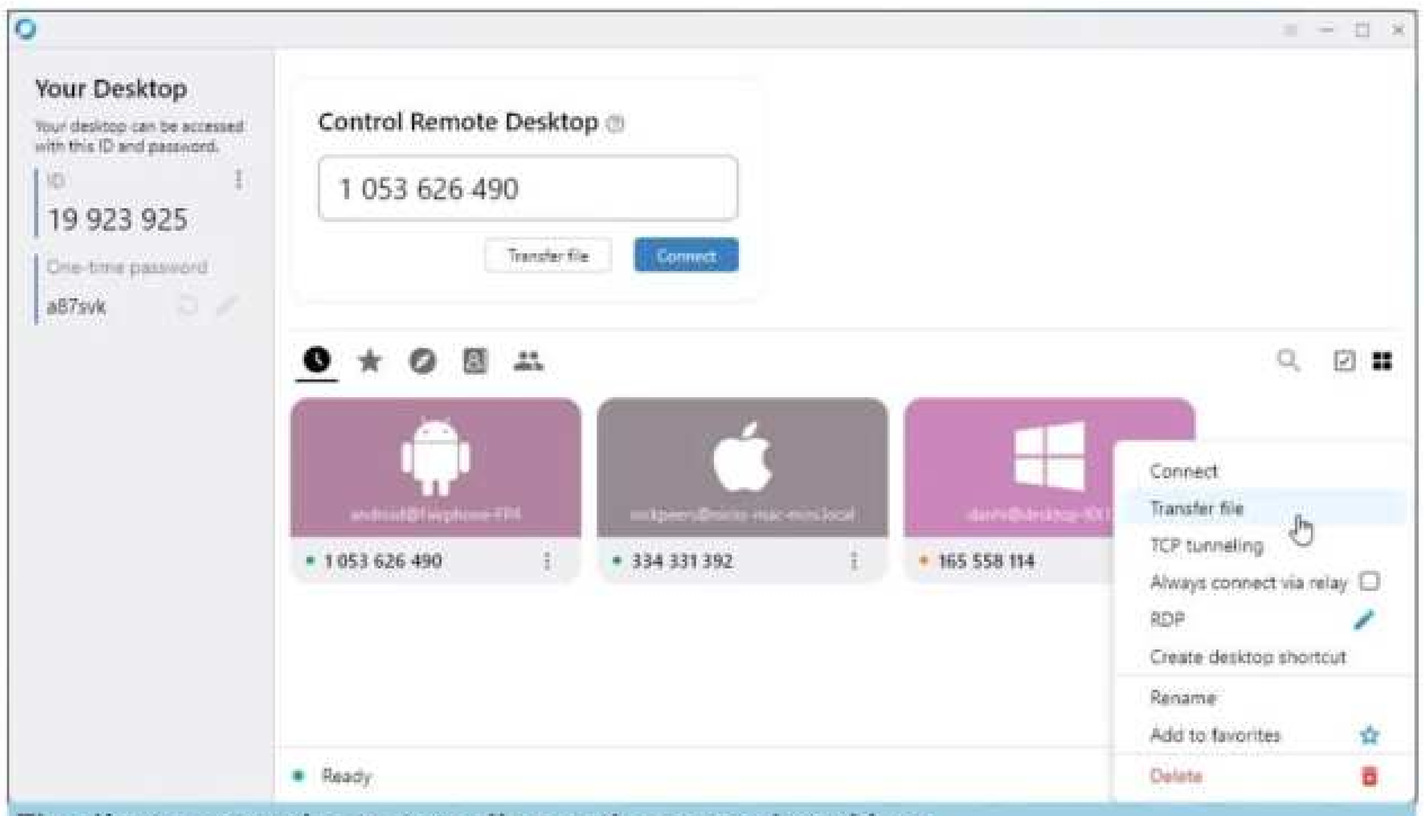

The client creates shortcuts to all recently accessed machines.

on using Command-key shortcuts over the regular Ctrl-key combos found in Windows, while 'Reverse mouse wheel' and 'Swap left-right mouse button' do exactly what they say.

MORE CLIENT TOOLS

There are five other buttons to be found on the pop-up menu—to keep it permanently visible, click the pin button on the left. Click the lightning button next to it, and you'll see various system-related options. These include an option for storing the remote device's OS password, enabling you to quickly enter it with a single click when the prompt appears on the remote machine (you'll find this particularly useful when you're connected to a Mac or Linux computer).

You'll also see a 'Transfer file' option, which opens RustDesk's file transfer tool. This requires an additional prompt for your connection password (or agreement from the person at the other end of the connection) before the main window pops up. You'll see the root directory on both machines, enabling you to navigate to any folder and easily send files between them.

Another interesting option is 'TCP tunneling', which is a means of setting up port forwarding between your device and the one to which you're trying to connect: Input your local port, the remote host address, and the remote port to which you wish to direct traffic. Once configured, you can connect to the remote device in different ways—for example, you could set a port forward to 22 on another device on your network using its IP address to allow you to connect via SSH using a terminal or command prompt window.

Here are three more options: 'Restart remote device' does what it says on the tin. You'll lose the connection, so wait a few minutes for it to restart before attempting to connect again. 'Insert lock' locks the remote screen should you need to do so for security reasons (don't forget to set up the 'OS password' option to speed up unlocking). Finally, 'Refresh' can be used to update the display.

When helping friends and family with their computer woes, being able to communicate what you're doing and why is a key part of any assistance process. Naturally, RustDesk has tools on hand to help—you'll find text chat and voice call options when you click the speech bubble button. Text chats appear in a floating pop-up window on your machine—you can reposition this using drag and drop. The chat appears in the access permissions pop-up on the remote machine.

The final options you'll see are a blue Record button—selecting this records the screen as a video, which is saved to a RustDesk folder inside your personal Videos folder—and a red cross icon

MORE DISPLAY OPTIONS

RustDesk's Display menu gives you a host of options for controlling your remote display, some of which—adaptive versus scaled, changing resolution—we cover in the main text. Other options are concerned with balancing the image quality with the speed of your connection. Explore the Image Quality submenu to experiment with bitrate and FPS settings if you find you need a trade-off between image quality and responsiveness of the mouse and keyboard. You can see the current state of your connection by checking 'Show quality monitor' from the Display menu—this opens a draggable pop-up inside the RustDesk client window that displays stats including connection speed, FPS, delay, and target bitrate.

The Codec drop-down enables you to change video codec—recent updates have enabled hardware codec support, making x265 the new default on supported devices. Other options are x264 (hardware) or VP9, VP8, and AV1 (all software).

Other options on the Display menu include being able to add virtual displays to supported remote OSes (such as Windows). When added, you can arrange items across multiple displays and move between them via this menu. Beneath this are checkboxes for 'Show remote cursor' and 'Zoom cursor'—by default, your remote device's cursor is hidden in favor of your own, so check the first option to track the remote cursor when it's moved by the device owner. You'll also find checkboxes enabling you to mute the remote desktop's sound, disable its clipboard, and lock the remote desktop after you end the session—for security reasons, obviously. Some devices also support Privacy Mode—when enabled, the device's display is hidden so no one can see what's happening unless they press Ctrl-P (make sure you block user input to circumvent this if you're accessing one of your own devices remotely).

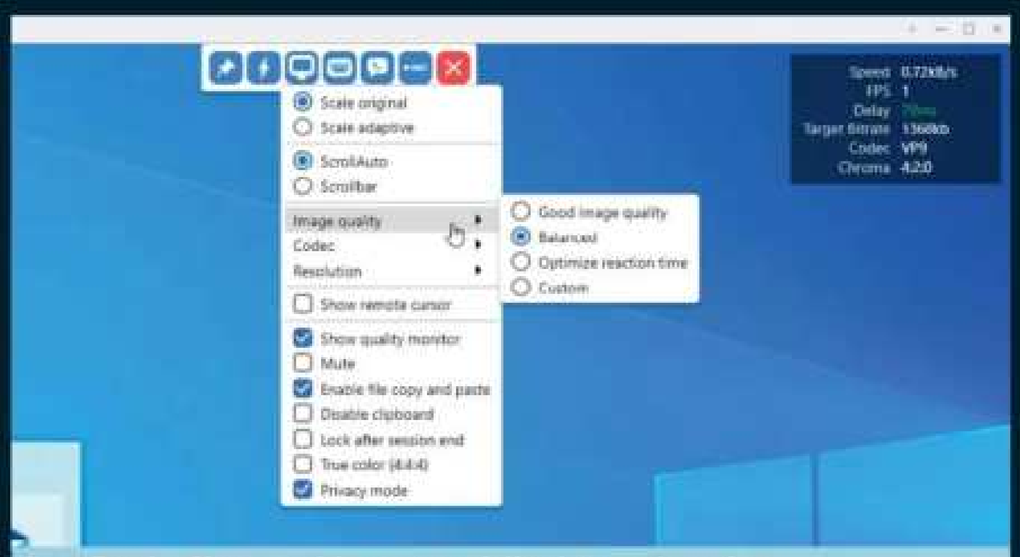

RustDesk offers a wide range of display options for your remote connection.

for providing an alternate way to end the session.

CONFIGURE CLIENT SETTINGS

After ending your connection, you'll find the RustDesk client has conveniently stored a shortcut to the remote machine in its History section for future use. This displays an icon indicating the remote device's OS type, the user and device name, and a handy colored dot that's green when the machine is online and accessible. You'll also see a vertical ellipsis button—click this to access all connection types, including 'Transfer file' if that's all you want to do, plus an option to rename the connection to something more descriptive. You can also add it to your favorites or delete the connection.

Other program settings—and there are many—are found by clicking the hamburger icon at the top to open the Settings tab. Thankfully, most settings are self-explanatory and enable you to set default behaviors for the various options we've explored. One section you shouldn't skip over is Security—select it in the left-hand pane and then click 'Unlock security settings' before clicking 'Yes' to give RustDesk the additional access it needs.

Permissions enables you to pre-select your choice of permissions when you—or someone else—connects to this device. The Password section is where you can change how connections are made to your current device—you can force remote users to always enter a password or require permission from your end.

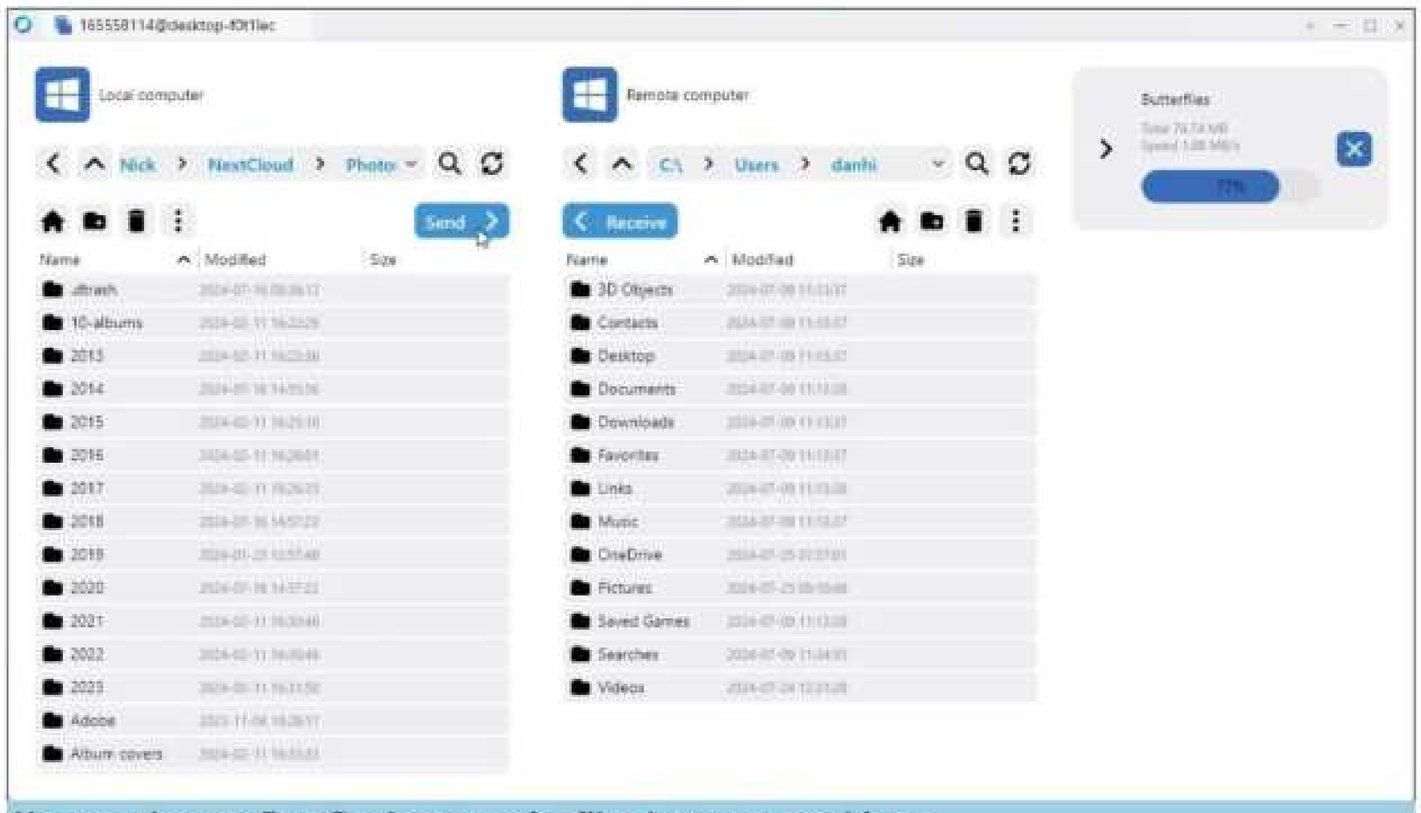

You can also use RustDesk to transfer files between machines.

By default, RustDesk always assigns a new one-time password each time you launch the client—it's six digits in length, but you can increase this to eight or 10 digits, or replace it with a permanent password if you plan to use RustDesk for remote connections to this machine. Below this you'll also find a 2FA option, which is only available when connected through the public RustDesk server, as well as an option to change the ID.

There are a host of additional options below this, but one stands out: 'Enable direct IP access'. If you'd like to connect directly to another device on your network using its local IP address—thus bypassing the RustDesk relay server to speed up performance—check this option. Once checked, a remote device can be accessed simply by inputting its local IP address (typically 192.168.x.y) into the 'Control Remote Desktop' box instead of its ID.

SET UP YOUR OWN SERVER

RustDesk's public servers do a good job but can be sluggish when lots of people are using them. By setting up your own self-hosted server, you can speed things up. RustDesk requires you to install two server components that manage the

MAKE YOUR RUSTDESK SERVER REACHABLE

Anyone wishing to connect through your RustDesk server needs a permanent address to point at. This could be your public IPv4 address (visit http://whatismyip.com to find out what this is), but apart from being difficult to remember, it's also likely that the address changes over time—for example, when you reboot your router. In previous issues we've talked about setting up a subdomain or dynamic domain (such as those offered for free at www.noip.com) to redirect to your public IPv4 address, updating as required with the help of a DDNS client such as ddclient (https://ddclient.net) or one supplied by your domain or dynamic DNS host provider.

Because of the way RustDesk communicates—through its own unique set of ports—you don't need to set up a dedicated subdomain. If you've followed one of our tutorials on running self-hosted services such as Vaultwarden, Jellyfin or your own WireGuard VPN server from the June 2024 issue, you may have a subdomain or dynamic DNS domain in place that you can use. Make sure you choose a subdomain that's protected with an SSL certificate—we recommend Nginx Proxy Manager (https://nginxproxymanager.com) for a simple way to obtain this if you've not already done so. Visit https://linuxformat.com/archives?listpdfs=1 and click the 'Nginx proxy' link for a comprehensive guide to setting up and using it.

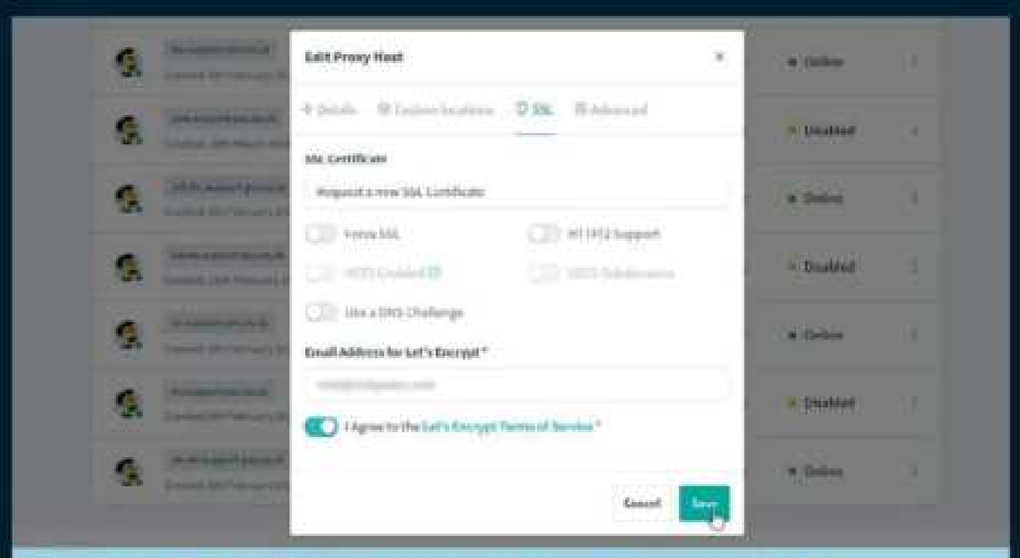

Use Nginx Proxy Manager to equip your chosen domain with an SSL certificate.

Once your subdomain and SSL certificate are in place, you can then input the domain into the 'ID server' and 'Relay server' fields of the ID/Relay server section under RustDesk's 'Settings > Network' dialog. And once connected, verify it's an encrypted connection by looking for the green shield icon in the RustDesk client window. Hover your mouse over this and it should confirm the connection is secure with a 'relayed and encrypted' status message.

connection between devices. They can be installed on your main Windows PC, or any server or NAS capable of running Docker or Podman containers, including QNAP, Synology, and our small NAS build from the April 2024 issue.

Before you begin, you need to open several ports to allow RustDesk to work remotely, so log on to your router and access its port forwarding section, then forward the following ports:

TCP: 21115, 21116, 21117
UDP: 21116

These are the only ports you need to open if you stick to the free OSS implementation of RustDesk—if you purchase a Pro license (see https://rustdesk.com/pricing.html—it's pricey at $9.90 a month), you also need to open TCP ports 21114, 21118 and 21119 to gain access to the web console and to support web-based connections.

These ports need to be forwarded directly to the IP address of the PC your RustDesk server is installed on, which means it needs a static IP address. If you install it on a server following the guide in the box to setting up a Podman or Docker container, this should already be the case, but if you're using your own Windows 11 PC, you may need to configure this.

You have two choices: the first is to explore your router's configuration utility for a DHCP Reservation section, which instructs the router to always assign the same IP address to specific devices on your network. Alternately, configure your PC to set its own static IP address. To do the latter, navigate to 'Settings > Network & Internet > Ethernet' and click 'Edit' next to 'IP assignment'. Select 'Manual' from the drop-down menu, flick the IPv4 switch to On and fill in the details: Choose an IP address that isn't in use elsewhere on your system or part of your router's reserved DHCP address space (use a tool such as Advanced IP Scanner from

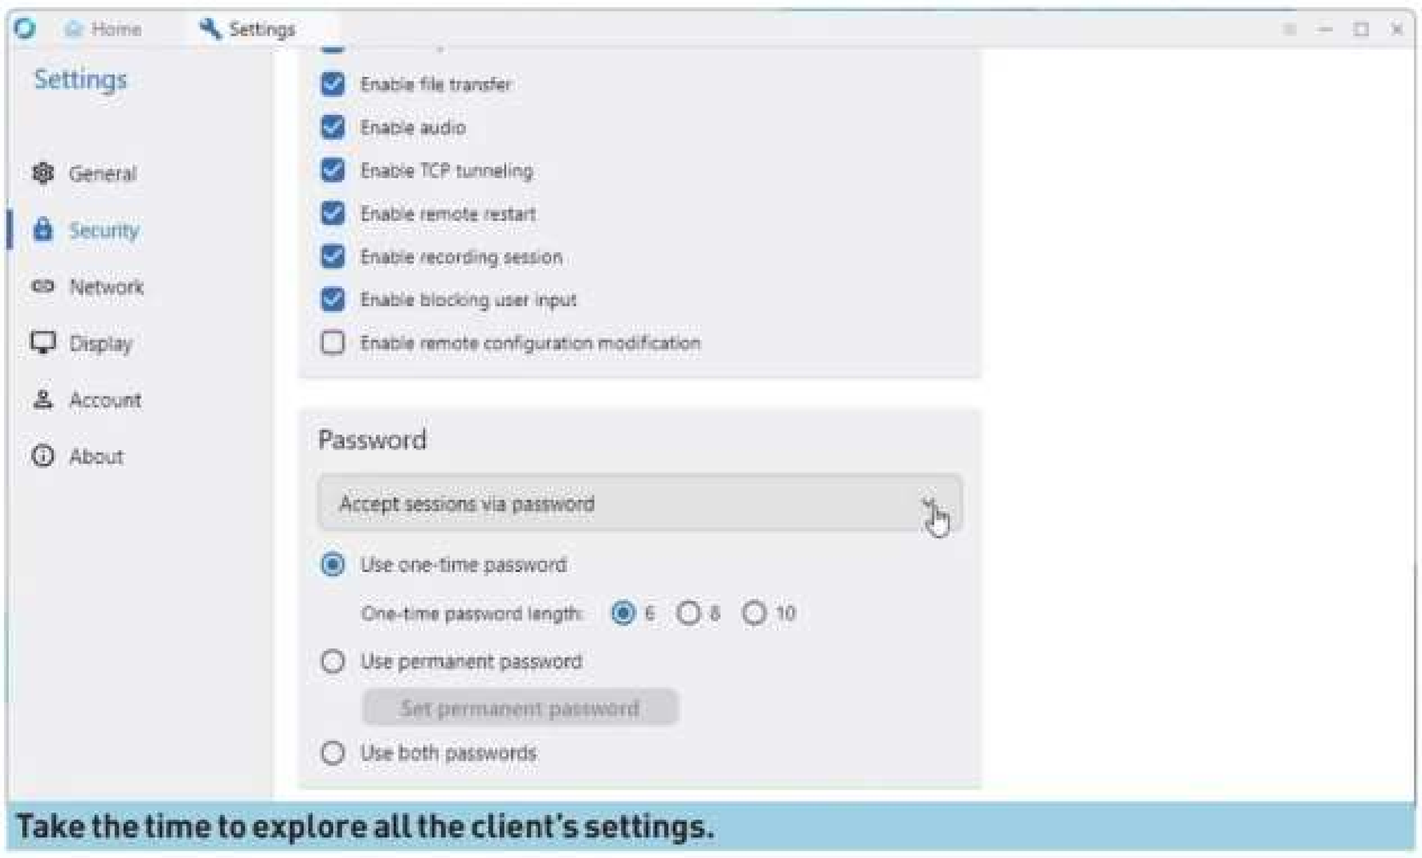

Take the time to explore all the client's settings.

www.advanced-ip-scanner.com if you're unsure what IP addresses are free to use).

Now fill in the other details: Subnet mask is usually 255.255.255.0, while the Gateway should be your router's IP address. When it comes to the preferred DNS server, you have up to three choices: your router's IP address, 1.1.1.1 (for Cloudflare) or 8.8.8.8 (for Google), or the IP address of your Pi-hole server if you've set one up (see the January 2024 issue).

Finally, anyone connecting to your RustDesk server (including you when away from home) needs to know your home network's public IP address, which is assigned to you by your Internet provider, and is that shown at www.whatismyip.com when you visit the website at home. If this is a static IP address, you can move on to the next step; if not, you might want to consider connecting through a subdomain or dynamic domain. See the box on the previous page for instructions.

INSTALL WINDOWS SERVER

You're now ready to install the RustDesk server. If you plan to install it on a Windows

Input the contents of 'id_ed25519.pub' into the 'Key' field.

computer, you have a choice of two approaches, both of which are described in detail at https://rustdesk.com/docs/en/ (expand 'Self-host' and then click on 'Windows & PM2 or NSSM'). We're going with the PM2 version, so the server will only run when your PC is switched on and you're logged into your own user account. That's a logical assumption if you want to access your PC remotely or use RustDesk solely to provide remote assistance to other people.

The PM2 version requires you first install Node.js, an open-source JavaScript runtime environment. We recommend downloading the latest LTS version via the download link at https://nodejs.org/en—during installation leave the box unchecked when prompted to automatically install additional tools for native modules.

Once installed, it's time to install PM2. Press Win-R, type 'cmd.exe', and press Enter to open a regular command prompt (administrative access is not required or desired in this case). Now type the following command:

```
npm install -g pm2
```

When it's finished adding packages, type the following two lines:

```
npm install pm2-windows-startup -g
pm2-startup install
```

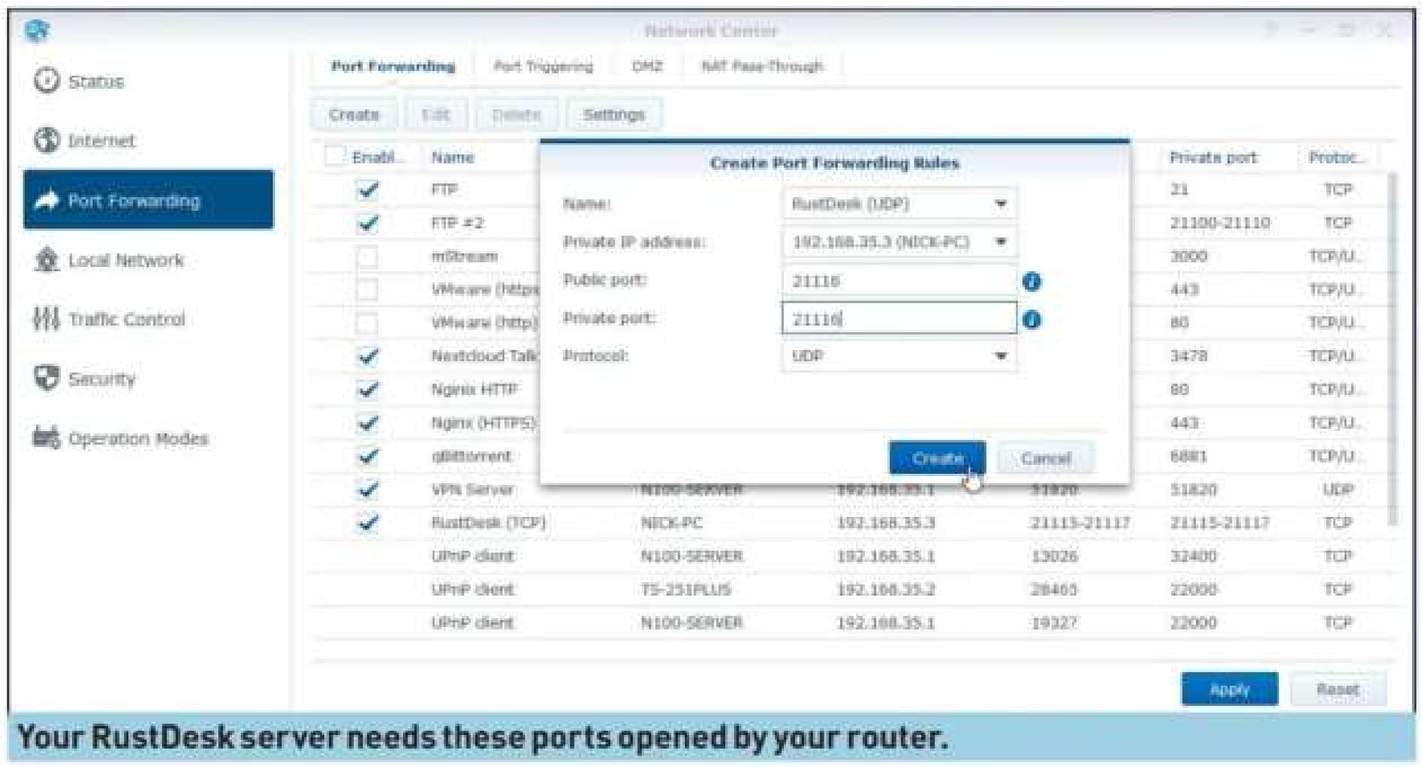

Your RustDesk server needs these ports opened by your router.

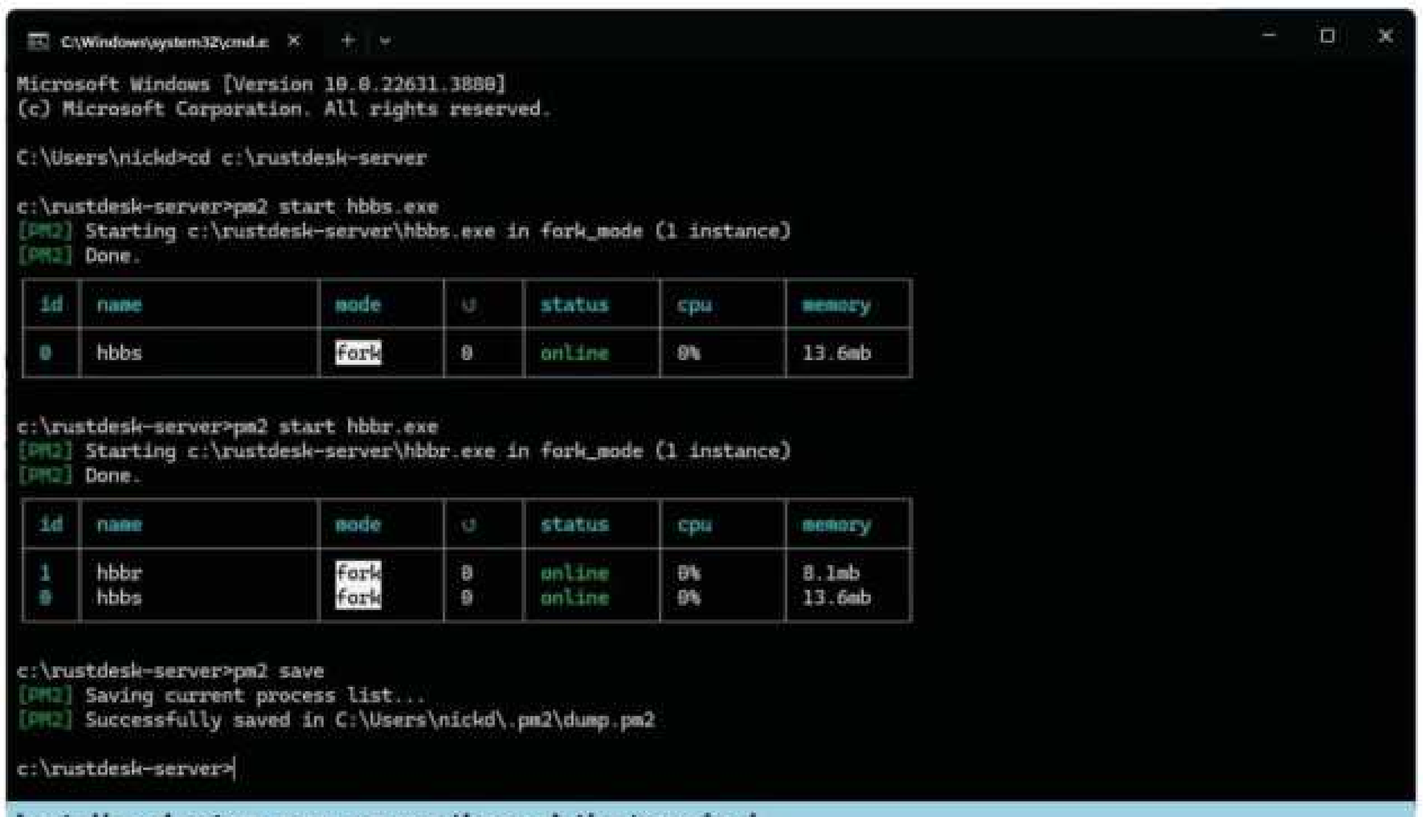

Install and set up your server through the terminal.

Now visit https://github.com/rustdesk/rustdesk-server/releases and expand all available releases under 'Assets' for the latest build, then double-click the 'rustdesk-server-windows-x86_64-unsigned.zip' link to save it to your PC. Right-click the zip file and choose 'Extract All'. When complete, open the 'rustdesk-server-windows-x86_64-unsigned' folder and rename the 'x86_64' folder inside to 'rustdesk-server' before copying it to the root of your C drive. Return to the command prompt window and issue the following commands to complete setup:

```
cd C:\rustdesk-server
pm2 start hbbs.exe
pm2 start hbbr.exe
pm2 save
```

CONFIGURE CLIENT

Go into the 'C:\rustdesk-server, folder, where you'll see two files of interest: 'id_ed25519' and 'id_ed25519.pub'. These are the private and public keys required for a secure connection to the server. Open the latter ('id_ed25519.pub') in your text editor and copy the contents to your clipboard.

Now you need to instruct each client to use your server instead of RustDesk's public servers. To do this, open the RustDesk client's Settings and select 'Network', then click 'Unlock network settings', then 'Yes'. All the configuration needs to be done in the 'ID/Relay server' section. First, paste the contents of 'id_ed25519.pub' into the 'Key' field, then input your chosen subdomain or static public IP address into both 'ID server' and 'Relay server' fields. Click 'Apply'.

Repeat the process on another client for testing purposes—how you securely transfer the public key is up to you, but we used our self-hosted Bitwarden server's secure notes feature to both store and copy the key easily between devices. Once configured, attempt to connect one of the configured clients to the other—all being well, the connection should work.

Once you've confirmed the connection, return to 'Settings > Network' and click the 'Export' button next to 'ID/Relay server'. If you have to use the public servers again for any reason, you can simply clear all three fields and click 'Apply'. When you want to reconnect to your own server, return here and click the 'Import' button to restore the server settings with a single click. Future versions of the client (1.3.x) will enable you to store multiple server configurations.

Congratulations, you've not only set up a robust, powerful (and free) remote access tool, but you've also added a layer of security and potentially improved performance by running it through your own RustDesk server, too. All that's left to do is configure all other clients with your server details, and you're good to go.

INSTALL IN DOCKER/PODMAN

If you're looking to provide a RustDesk server with 24/7 access, your Windows PC may not be the best choice to host it. The good news is that if you have a suitable server running Docker or Podman, you can get your RustDesk server up and running with the minimum of fuss. Example scripts for Docker users can be found at https://rustdesk.com/docs/en/self-host/ (select 'Docker' under '2.1 RustDesk Server OSS'), so we're focusing on how you'd adapt those scripts for a rootless Podman installation, as we've outlined in recent issues.

First, obtain the latest server image:

```
podman image pull docker.io/rustdesk/rustdesk-server
```

Next, create a suitable folder inside which you'll store your RustDesk configuration, including the all-important key—for example, the following line would create a hidden 'rustdesk' folder inside a 'container-data' folder within your home folder:

```
mkdir ~/container-data/.rustdesk
```

Then consult the screenshot on the right, where you'll see the commands you need to enter to create a pod containing both the main and regular server containers. You need to adapt the lines marked '-v' to point to wherever you placed the hidden 'rustdesk' config folder, plus make sure 'hbbs -r' is pointed to your own subdomain. Once both containers are up and running, navigate to the hidden '.rustdesk' folder where you'll find the private key you need in the 'id_ed25519.pub' file—open it using a text editor such as Nano, and then copy and paste the key somewhere suitable before passing it on to all the devices that need to connect through your server.

Once you've confirmed the connection works as it should, you can then complete the process by configuring RustDesk to automatically start with your server—again, the screenshot reveals the commands to do so on your Linux-based server.

```
podman pod create --name=rustdesk --
net=host

podman run -td \
  --pod=rustdesk \
  --name hbbs \
  -v ~/container-data/.rustdesk:/root
\
  docker.io/rustdesk/rustdesk-server \
  hbbs -r your.domain.com

podman run -td \
  --name hbbr \
  --pod=rustdesk \
  -v ~/container-data/.rustdesk:/root
\
  docker.io/rustdesk/rustdesk-server \
  hbbr
```

TO ENSURE RUSTDESK AUTO-STARTS WITH YOUR SERVER:

```
podman generate systemd --new --name
rustdesk -f
mv -v pod-rustdesk.service
~/.config/systemd/user/
mv -v container-hbbs.service
~/.config/systemd/user/
mv -v container-hbbr.service
~/.config/systemd/user/
systemctl --user daemon-reload
systemctl --user enable pod-
rustdesk.service
systemctl --user enable container-hbbr
container-hbbs
```

These commands should get RustDesk server running in Podman.

CENTERF

Dawn of a new era

IT'S FINALLY HERE. Notebooks, laptops, ultrabooks—all of them have led to this point in time. We're finally witness to an immeasurable change in how our portable world of computing will be moving forward.

The world of TN, IPS, and VA screens is coming to an end, and with it the advent of AI. We're now at a point where OLEDs are so cheap, and AI processors too, that it's hard to see a world where the two don't exist in every portable powerhouse. Arm, Ultra Cores, Ryzen 8000 series—all feature AI, and all typically come with that OLED panel. That crisp screen, with rich color accuracy, insane frequencies, and rapid response times.

Take the Asus Zenbook Duo we have here. By now, you'll likely have read our words reviewing the titular halo product, but it encapsulates where the industry is headed.

In some ways, there's a sadness to the lack of new technology or unique hardware perspectives. But in others, it means more time and resources invested into pushing them above and beyond what they are today. Yes, it seems we have arrived at a precipice. The question is, will x86 last? It's time to turn the page. **–MAXIMUM PC**

1 ASUS ZENBOOK DUO

Potent, portable, powerful, and perpetually configurable. We took a look at the Asus Zenbook Duo in our May 2024 Edition, complete with its own Neural Processing Unit, and twin OLED panels. It's outstanding.

2 X86 OR ARM?

Arm is starting to make some serious headway against our oldest processing friend, the x86 architecture. With the Snapdragon X Elite now pulling ahead quite competitively against Intel and AMD's latest and greatest, it may only be a matter of time until the switch to RISC is here.

3 OLED SUPREMACY

OLED tech is thinner, lighter, more efficient, and easier to manufacture than traditional screens. Combine that with far greater color range, accuracy, brightness, and super-low response times and high refresh rates, and it has quickly become the ultimate panel tech.

2024

CPU Architecture DEEP DIVE

New AMD and Intel processors take center stage

by Jarred Walton

AMD AND INTEL have traditionally released new processors every year, often with the biggest updates arriving for the fall season. 2024 continues that trend, but there are new processors, and there are *new* processors. In what has become mostly a two-year cadence, both AMD and Intel have major architectural overhauls going on, promising bigger changes than the off-year refreshes.

It's not just about desktop CPUs, either. As another sign of the current state of affairs, where mobile solutions continue to outsell desktop parts, both companies have laptop-centric designs coming out. Intel even leads off with its mobile-focused Lunar Lake architecture, which has some important differences from the upcoming desktop-centric Arrow Lake designs that we'll cover here. AMD likewise has its Granite Ridge line of CPUs that target desktops, with Strix Point processors going after the mobile market. Both use AMD's latest Zen 5 architecture, but there are plenty of under-the-hood differences to discuss. There are server designs from both companies happening as well, but we won't get into those.

All told, we have a collection of CPU architectures to discuss, with numerous changes from the traditional approach we've seen in the past. Intel is moving away from being the sole or even primary manufacturer of the silicon used for many elements of its latest CPUs (see last month's Tech Talk, MPC233), turning to rival TSMC for many aspects of its multi-tile solutions. AMD has used TSMC for years, and will continue to do so for this round of updates, but there's news to discuss there as well.

There's a lot to cover, so turn the page to join us as we seek AMD's Zen for the fifth time while swimming in some Intel lakes.

AMD Zen 5 Architecture

AMD IS FIRST out of the gate with its new Zen 5 microarchitecture, powering the latest Ryzen 9000-series desktop CPUs, as well as the Ryzen AI 300-series laptop parts. Officially, AMD has announced four desktop chips and three laptop parts, but we can expect that number to increase over time. Zen 5 gets further split into 'regular' Zen 5 cores, as well as density-optimized Zen 5c cores (see Zen 5c sidebar), but most of the functional units are the same.

Every new CPU architecture looks for ways to improve overall performance, using some combination of improved IPC (Instructions Per Cycle) and higher clock speeds. Adding performance alone isn't enough, however, as we learned in the

The block diagram overview of Zen 5 shows significant changes compared to prior Zen designs, with larger buffers and a wider pipeline.

© AMD

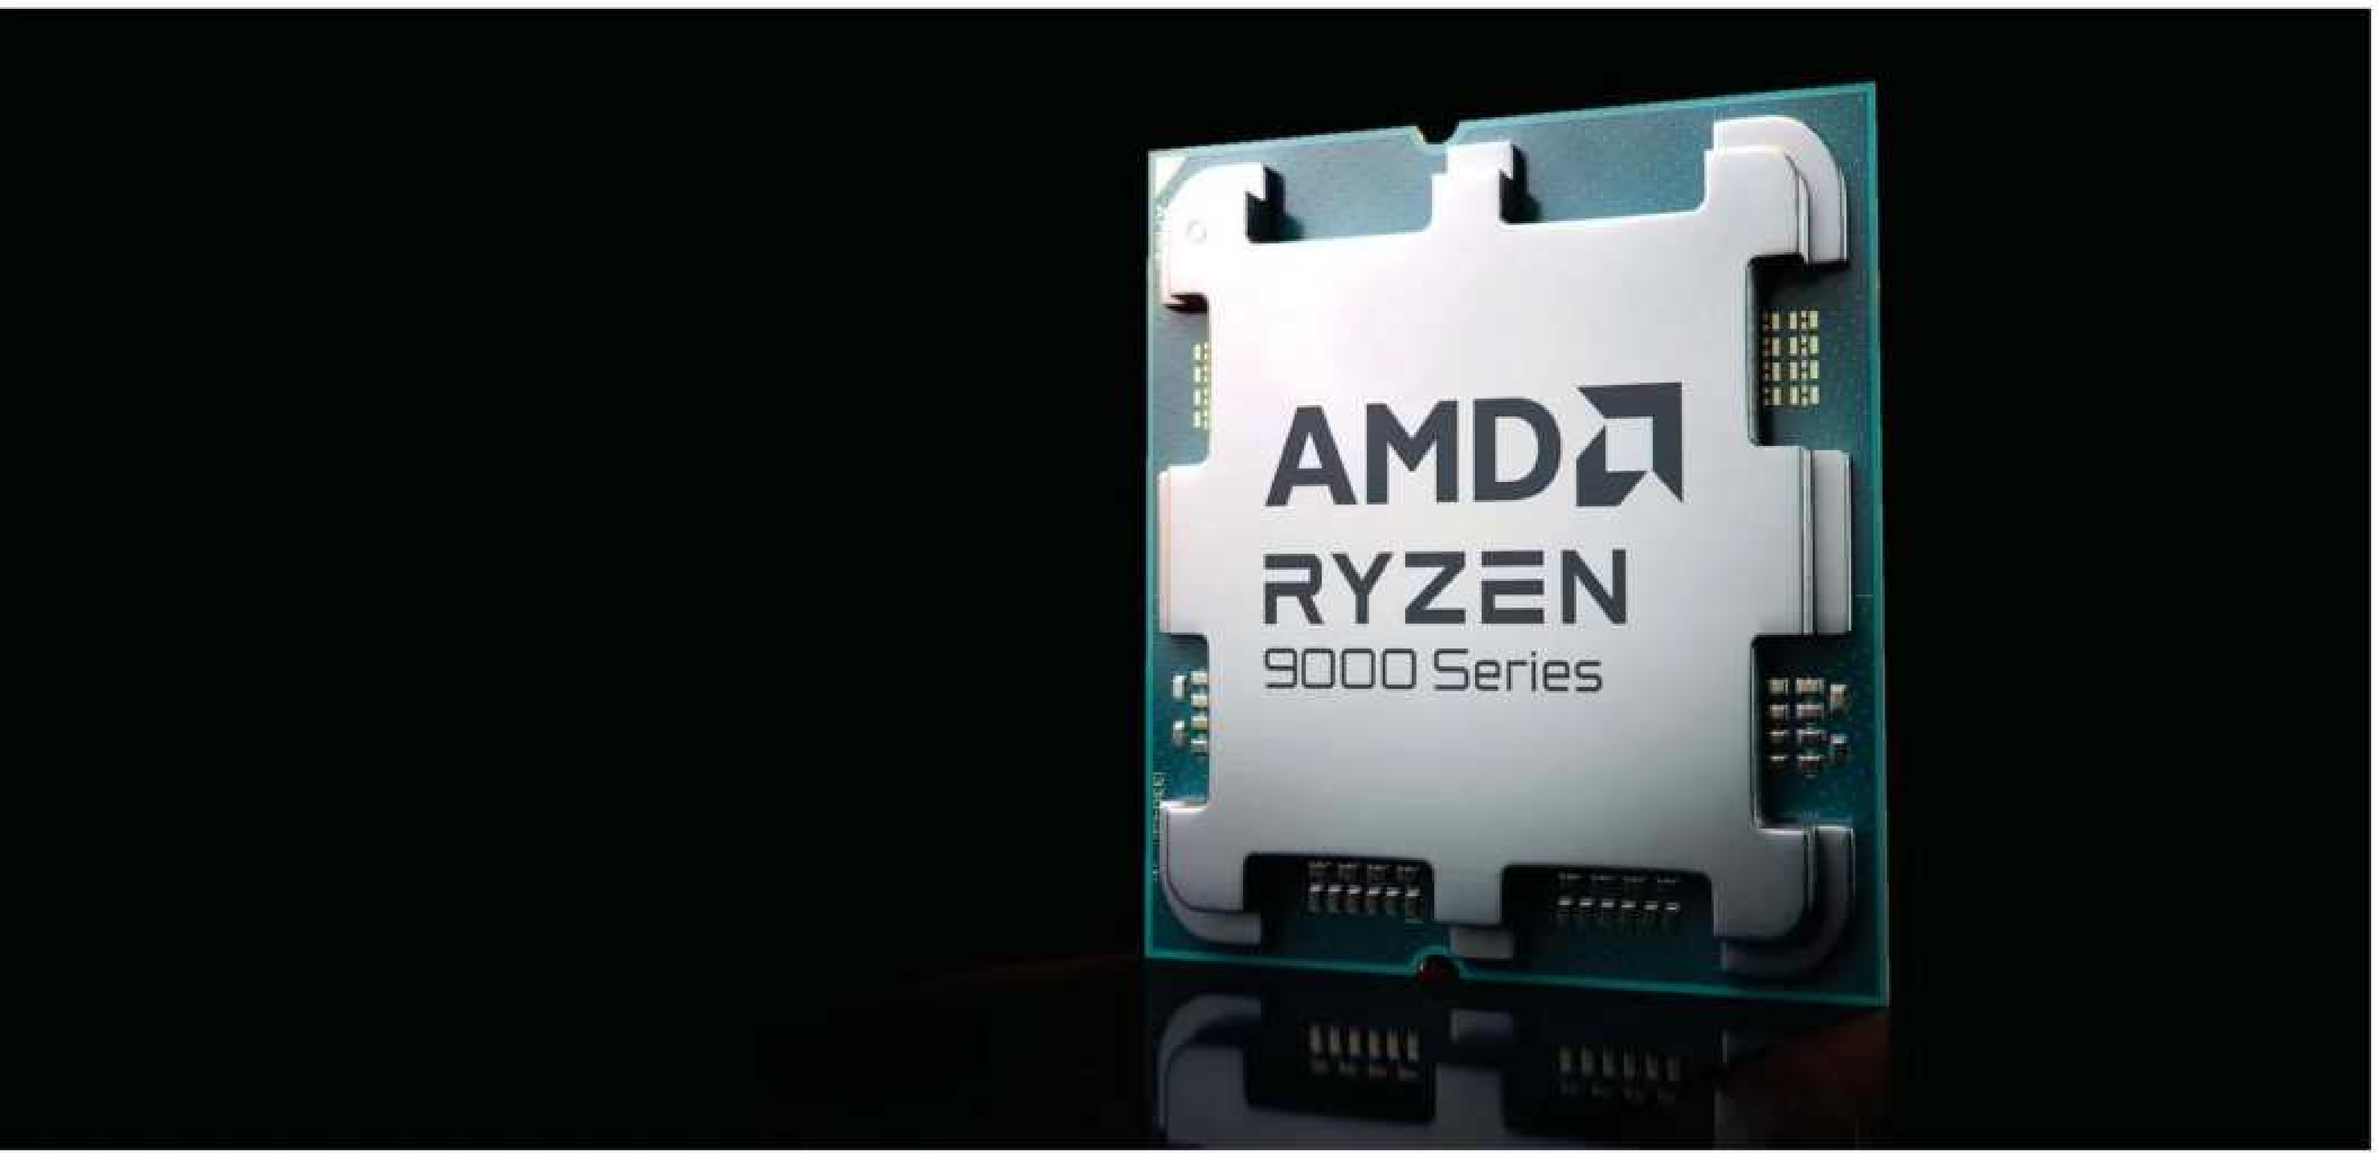

clock speed wars of the mid-2000s. Ideally, new architectures offer improved performance along with better efficiency, which is why clock speeds haven't seen massive changes over the past couple of decades.

The first CPUs to hit 1GHz came out in March 2000 with AMD's Athlon 1000 and Intel's 1GHz Pentium III models. Clock speeds doubled to 2GHz and more just 18 months later with Intel's Pentium 4 processors, and by late 2002 there were 3 GHz Pentium 4 chips available. That's when CPU designers ran into a wall and clock speeds began increasing much more slowly. We didn't get a retail 4GHz processor until ten years later with the AMD FX-4170! AMD's infamous FX-9590 broke the 5GHz mark in 2013, but with rather extreme power draw for the time, and we only just hit 6GHz CPUs with the Core i9-13900KS in early 2023.

> Every new CPU architecture looks for ways to improve overall performance, using some combination of improved IPC (Instructions Per Cycle) and higher clock speeds

What we're saying is that increasing CPU clocks generally hasn't been the best way to boost performance for over two decades, so instead, CPU designers have focused on boosting IPC. There are multiple tactics that can improve IPC, including wider data paths, deeper buffers, more execution units, better branch prediction, simultaneous multi-threading (SMT), and larger caches. Zen 5 basically ticks off every one of those boxes in its quest for higher instruction throughput.

AMD reworked the front end for Zen 5 to include dual 4-wide instruction decoders, along with an 8-wide dispatch unit. Previous Zen architectures were 6-wide designs, so this represents a potential 33 percent increase in instruction throughput. The branch prediction has also been improved to have lower latency and improved accuracy—key factors for increasing performance.

L1 cache throughput has been improved, and the L1 data cache has increased to 48KiB, up from 32KiB on prior Zen chips. Along with the wider dispatch unit, there are more execution units. For integer work, there are six ALUs (Arithmetic Logic Units), three of which support multiplication, alongside four AGUs (Address Generation Units). Floating-point and vector units support full AVX-512, with six pipelines to handle the various instructions.

Together, the above changes provide an average improvement of around 16 percent relative to Zen 4, according to AMD. That's a significant jump, and represents a continuation of AMD's claimed double-digit percentage gains for each Zen generation. Zen 1 offered a 52 percent increase in IPC compared to the previous Bulldozer architecture, Zen 2 was a 15 percent improvement; Zen 3 was 19 percent, and Zen 4 was 16 percent. It's a great track record for AMD's CPU team.

Zen 5 desktop CPUs continue AMD's chiplet tradition, with a single IOD (Input Output Die) that has the memory controllers, basic graphics support (two RDNA 2 CUs), PCIe lanes, USB controllers, and more. It's the same IOD as with Zen 4, built on TSMC's N6 node to save on costs, so base connectivity will be the same as Zen 4. The CCD (Core Chiplet Die), meanwhile, uses TSMC's N4P node, a refinement of N5 that improves performance and efficiency. AMD will also offer Zen 5 CCDs that use the newer N3B node, though that will presumably be for EPYC variants, rather than Ryzen CPUs.

AMD Strix Point

FOR LAPTOPS, AMD's Strix Point makes similarly major changes to the underlying architecture for CPUs that will be sold under the Ryzen AI brand. As with the previous-generation Phoenix Point and Hawk Point, this is a monolithic design—no chiplets here. The single large die contains all the IO connectivity for memory and PCIe devices, CPU cores, an RDNA 3.5 GPU, and an XDNA 2.0 NPU (Neural Processing Unit).

The top two Ryzen AI models (Ryzen AI 9 HX 375 and 370) have 12 CPU cores: four standard Zen 5 cores and eight Zen 5c cores. All the cores have SMT enabled, letting the chips execute 24 concurrent threads of work. The previous generation 7000- and 8000-series laptop CPUs—which were basically the same architecture—topped out at eight cores: four Zen 4 and four Zen 4c, so in addition to architectural advances, the Ryzen AI processors have a 50 percent increase in core count.

As impressive as the CPU aspect might seem, it's the other areas of Strix Point that show real promise.

AMD's Ryzen AI processors use a large 232.5mm2 die built using TSMC's N4P node—that's 68 percent larger than Hawk Point's 138mm2 die.

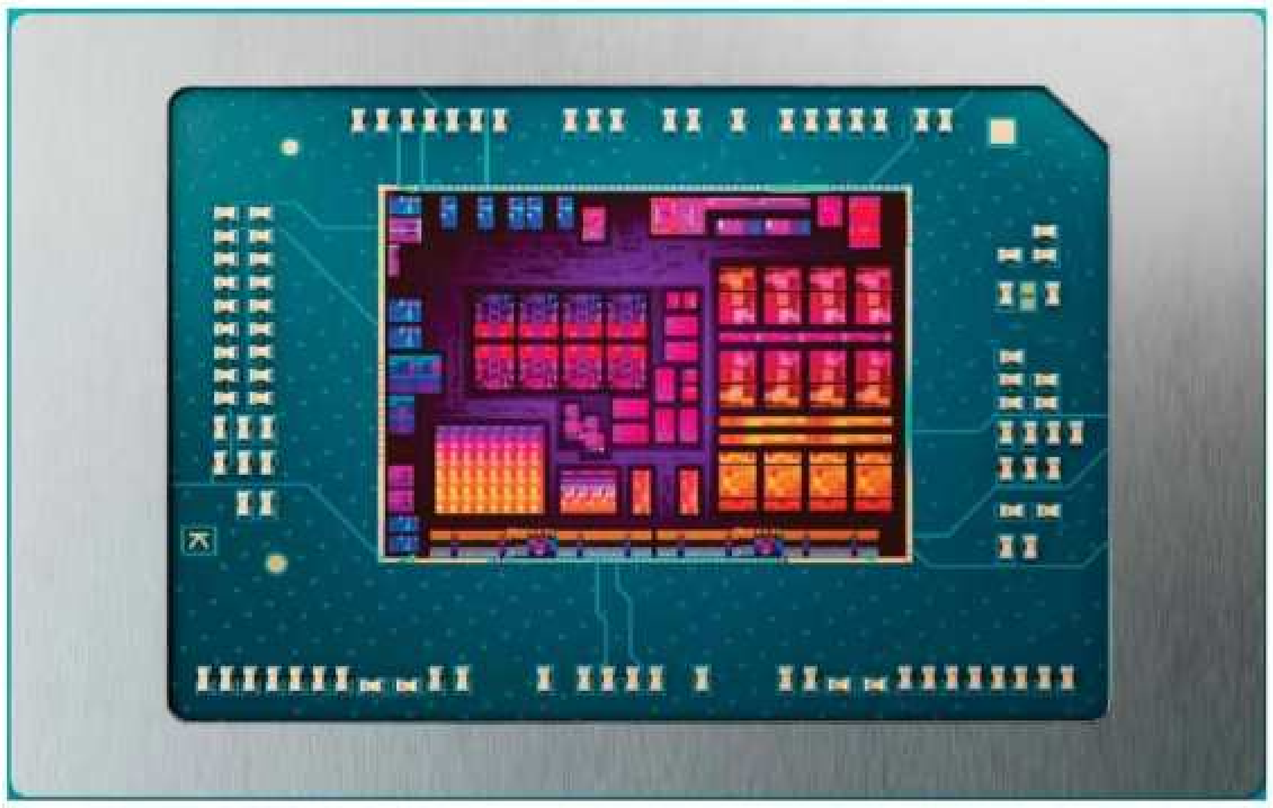

Hawk Point and Phoenix Point offered integrated RDNA 2 graphics with up to 12 CUs (Compute Units), each with 64 shader ALUs. Strix Point moves to the RDNA 3.5 architecture, and then increases the CU count to 16. That's at least a 33 percent potential boost to graphics performance, and likely more than that for real-world workloads.

As with the CPU architecture, it's not just about adding more cores and improving performance. RDNA 3.5 contains optimizations specifically to boost efficiency, with double the texture sampling rate, twice the interpolation and comparison rates for shaders, and improved memory management, with better compression techniques and additional tuning for LPDDR5 memory.

The third leg of Strix Point is the new XDNA 2 NPU, which support both INT8 and FP16 data types to better handle a variety of AI workloads, along with a new 9-bit Block FP16 data type that offers nearly the precision of FP16 with close to the performance of INT8. AMD's Phoenix Point processors (7040-series) were the first to include an NPU, rated at just 10 TOPS (INT8 teraops); Hawk Point was basically the same chip, but with a boost in clocks allowing up to 16 TOPS of performance. Strix Point's NPU sees an expansion from 20 AI tiles to 32, with each tile offering much higher performance, yielding a total of 50 TOPS of compute.

The net result of all these elements is that not only do the new Ryzen AI chips have higher CPU, graphics, and AI performance than the previous generation parts, but the fastest models are also rated for 28W TDP compared to 45W on the previous top-end Ryzen 9 8945HS. Strix Point is also the first x86 chip that's Microsoft Copilot+ certified, although we're still waiting to see what that means in practice after Recall got sent back to the drawing board.

Intel Lunar Lake Architecture

INTEL HAS MULTIPLE NEW CPU architectures in the works, although the desktop and laptop parts appear to be even more wildly divergent than what AMD is doing with Zen 5. The first parts to launch will be the Lunar Lake processors for laptops and other portable devices. Much like last year's Meteor Lake, Intel is leading with mobile chips, building on many elements of the previous generation, but with newer, faster, and more efficient designs.

As a recap, the previous-generation Meteor Lake (Core Ultra) processors use a multi-tile approach with four tiles, plus the base Foveros tile. The CPU tile is made using the Intel 4 process node, and sports a hybrid architecture similar to Alder Lake and Raptor Lake desktop parts, with Redwood Cove P-cores and Crestmont E-cores. The GPU tile uses TSMC's N5 node and handles graphics work. The largest tile is the SoC tile made on TSMC's N6 node, and it includes an NPU, two lower power E-cores, the memory controllers, and video decoding support, plus display outputs. Finally, the IO tile also uses TSMC N6, with Thunderbolt 4 and PCIe 5.0 connectivity support.

Lunar Lake builds on Meteor Lake's concepts, but it condenses things into just two main tiles this time: a large Compute tile that's made using TSMC's cutting-edge N3B node, and a Platform Controller tile that's made on the older TSMC N6 node. The only Intel-made silicon for Lunar Lake is the Foveros base tile that links everything together, which is made on Intel's 22FFL (now called Intel 16), originally used to make 3rd Gen Core Ivy Bridge CPUs back in 2011. Moving to just two processing tiles should reduce costs somewhat, but outsourcing both of those to TSMC means that Intel will end up paying more money per chip as it only

> All of this is good, but one of the most noteworthy changes is that Intel has worked to make the core design far more portable between different process nodes.

© AMD

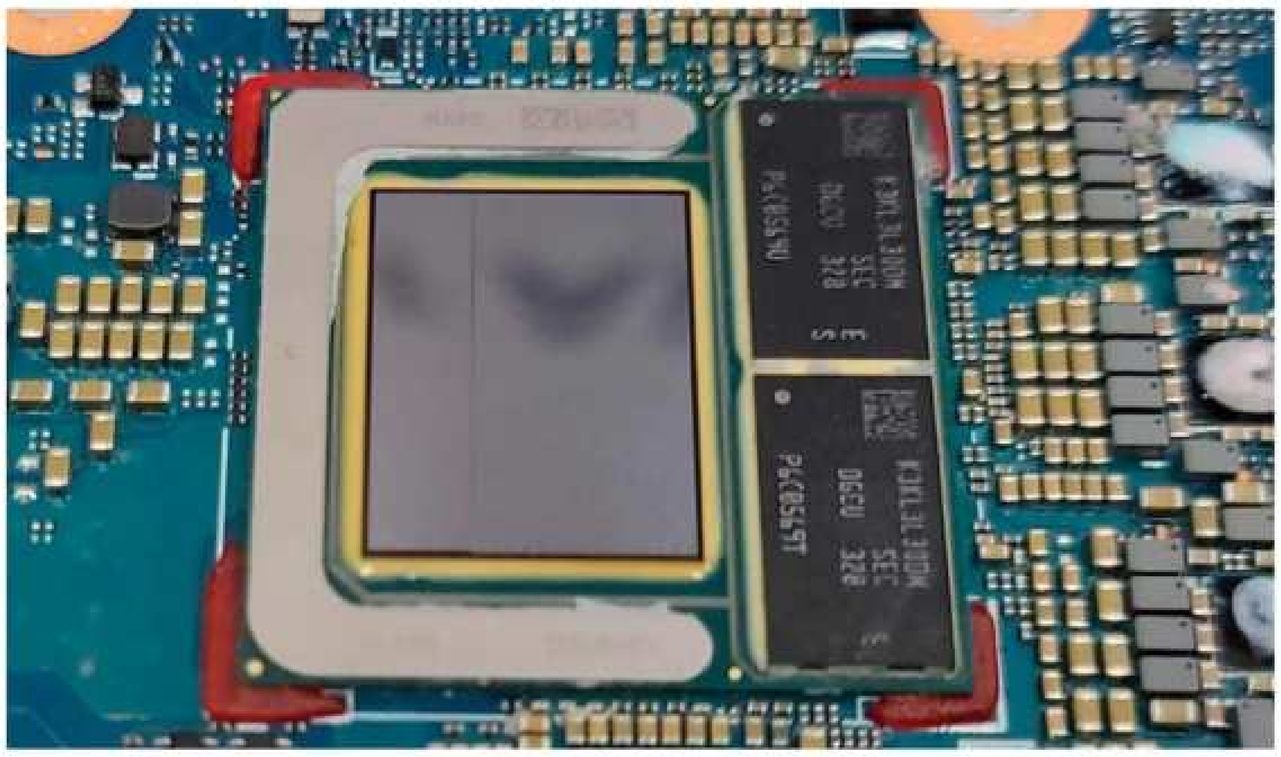

handles the packaging internally.

The Platform tile houses all the external I/O functions for the chip, including Wi-Fi and Bluetooth, USB 3.0 and 2.0, Thunderbolt 4, PCIe 4.0 and 5.0 interfaces, and the memory controllers. Most of these technologies don't need the latest process nodes, and so there are cost savings by using TSMC's N6 node. The Lunar Lake Compute tile is where all the exciting stuff resides, with four Lion Cove P-cores, four Skymont E-cores, a new Xe2 GPU, and the NPU 4.0 AI processor. We've got sidebars on the GPU and NPU, but let's tackle the two CPU architectures in more detail.

Intel's Lunar Lake package includes a large Compute tile, smaller Platform tile, and a tiny square 'filler' tile for stability, along with two LPDDR5x chips.

Intel's Lion Cove P-Cores

INTEL USES A HYBRID architecture on its processors, with performance-oriented P-cores and efficiency-focused E-cores. Lion Cove is the codename for the latest iteration of the P-core, following on from the Golden Cove architecture used in Alder Lake and Raptor Lake, though Intel renamed this to Raptor Cove for the 13th and 14th Gen chips.

A lot of the changes with Lion Cove echo what we saw with Zen 5. Intel is adding more execution resources, along with deeper buffers and other improvements. For example, Lion Cove uses an 8-wide design, up from 6-wide on the prior generation. It also can retire up to 12 instructions per cycle, up from eight. To help, it has 18 execution ports, up from 12 on Golden Cove.

All of this is good, but one of the most noteworthy changes is that Intel has worked to make the core design far more portable between different process nodes. Previous generations used large numbers of 'Fubs'—functional blocks—often with tens of thousands of cells each. These were designed for a specific node—for example, Raptor Lake and Alder Lake use the Intel 7 node. With Lion Cove, Intel is using larger partitions that can range from hundreds of thousands of cells up to millions.

Lion Cove is far more portable between different processes—for example, Lunar Lake will leverage TSMC's N3B node for the compute tile, while Arrow Lake will be the first chip to use the Intel 20A node. A key factor behind the change was Intel's stalling out on 14nm for so long, which allowed TSMC to pass it by while newer CPU architectures were delayed while Intel worked out the kinks in its 10nm process.

Both Arrow Lake and Lunar Lake will omit Hyper-Threading, but Intel has a separate variant of the Lion Cove P-cores where SMT remains enabled. That will be for P-core only designs used in upcoming Xeon 6 chips. Another interesting change is that Lion Cove will allow for CPU clock speed adjustments in 16.67 MHz steps as opposed to the 100MHz steps used on current CPUs.

Taken together, Intel says that Lion Cove will deliver an overall 14 percent average IPC improvement over the Redwood Cove cores used in Meteor Lake, which were basically the Intel 4 process node variant of the Raptor Cove and Golden Cove cores used in Raptor Lake and Alder Lake.

Intel's Skymont E-Cores

SKYMONT MARKS INTEL'S THIRD E-CORE design for its hybrid processors, following Gracemont in Alder Lake and Crestmont in Meteor Lake. While the 14 percent IPC improvement with Lion Cove sounds good, Intel claims up to a 38 percent improvement in IPC with Skymont on integer workloads, and up to

THE ZEN 5c 'COMPACT' CORE

As far as we know, Zen 5c won't make an appearance on the desktop parts, but will instead be used with mobile processors, as well in certain high core count EPYC server CPUs. Fundamentally, it has the same instruction set and features as the bigger cores, but AMD focused on reducing die size. The result is a core die size that's 25 percent smaller than the Zen 5 core, with reduced clocks and other tweaks to improve efficiency.

Zen 5 uses a standardized CCD for desktops and servers, with eight CPU cores and 32MB of L3 cache. It's broken down into two CCXs (Core Complexes), each with four cores and 16MB of cache. AMD can use one or two CCDs on a package for its desktop parts, and up to 12 CCDs on chips like EPYC 9684X—with 16 CCDs coming in future EPYC designs.

Zen 5c for laptops uses a dual-CCX design with two 4-core clusters and a combined 8MB of shared L3 cache. That's one fourth the L3 cache of a Zen 5 core. AMD has lowered the clock speed targets, using fewer but larger functional logic blocks that can be packed closer together. To put it another way, the last several hundred megahertz of speed on the big Zen 5 cores requires more space to dissipate heat and additional tuning of data pathways, resulting in each core being about 33 percent larger than Zen 5c.

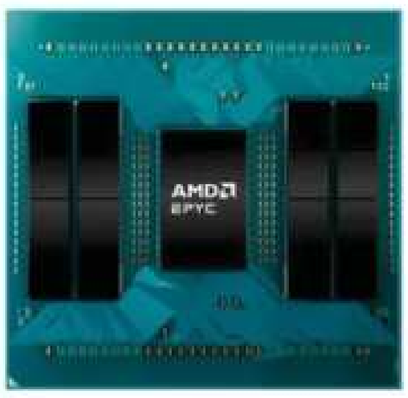

AMD's EPYC 'Turin' will use Zen 5c and pack 12 CPU cores per chiplet, giving it up to 192 cores.

Zen 5c still supports SMT (Symmetric Multi-Threading), and it's fundamentally different from Intel's P-core/E-core approach. Where Intel can pack four E-cores into the same area as a single P-core, AMD's Zen 5c dual-CCX (eight cores plus 8MB L3 cache) ends up being about the same size as a single Zen 5 CCX (four cores plus 16MB L3 cache).

Block diagram of Intel's Lion Cove performance cores that will be used in Lunar Lake and Arrow Lake CPUs.

a 68 percent improvement for floating-point work. That's huge, although there's a caveat that warrants mention: Intel is making those comparisons to the weaker low power E-cores used in Meteor Lake's SoC tile, not the standard quad-core cluster of E-cores used in the compute tile.

Lion Cove went wider with an 8-wide design, and Skymont takes that further with a 9-wide (3x3) decode cluster, up from a 2x3 decode on Meteor Lake. The allocation of instructions uses an 8-wide design, up from 6-wide in Crestmont. Skymont also has a larger out-of-order window, bigger physical register file, and deeper reservation stations for Int, Mem, and Vector work. The micro-op capacity is also 50 percent larger, with 96 entries (up from 64).

Intel targeted a doubling in vector performance, by including four 128-bit SIMD pipelines versus the two used in Crestmont. The E-core cluster has also received a bump in L2 cache capacity—up to 4MB from the 2MB previously offered.

It's a bit difficult to say precisely how much of an improvement Skymont will be, compared to Crestmont, but all indications are that even if the 38 percent IPC boost may be exaggerated, it's still likely to be in the 20 percent or higher range. Combined with the use of newer, more efficient process nodes—TSMC N3B for Lunar Lake and Intel 20A for Arrow Lake—Skymont should deliver some excellent generational gains in performance.

Intel Arrow Lake

ARROW LAKE CPUs will basically be the desktop alternative to the Lunar Lake mobile CPUs. After the relatively ho-hum Raptor Lake refresh last year with 14th Gen Core updates, Intel looks to deliver substantial changes and upgrades for desktop users. Intel hasn't officially confirmed specifications or model numbers, but the rumor mills and leak-mongers are running wild, and motherboards sporting the new LGA1800 socket were on display at this year's Computex show in Taiwan.

Like Lunar Lake, Intel will use a mix of Lion Cove P-cores and Skymont E-cores for Arrow Lake, which will also mark the switch to Intel's new Core Ultra branding and nomenclature. The top model is expected to have the same 125W TDP as the 14900K, but indications are that they'll be more efficient than the prior generation, and probably have lower peak power draw—the Raptor Lake 125W chips could draw up to 253W, and many motherboards would default to removing power limits, and could let the 14900K, as an example, draw over 350W! Arrow Lake CPUs should reign things back in, thankfully.

One of the biggest changes will be the lack of

BATTLEMAGE ARRIVES IN XE2 GRAPHICS

Xe2 is Intel's second-generation high-performance GPU, codenamed Battlemage—with Xe-LP, Xe-LPG, and Xe-HPG all counting as first-generation parts. We expect to see dedicated Arc GPUs using the Battlemage architecture before the end of the year, but the first Battlemage GPUs to hit the streets will come via Lunar Lake.

Intel's graphics division spent much of the past two years refining and improving its graphics drivers, and Xe2 will benefit from those efforts. In addition to all the efforts on the software and drivers side of things, Intel has repartitioned resources into a native SIMD16 engine to improve efficiency, added support for new data types, and boosted the ray tracing capabilities.

The Xe2 GPU in Lunar Lake comes with eight Xe-cores, the same count as in Meteor Lake, but under the hood there are significant differences. Xe-LPG in Meteor Lake had 16 256-bit vector engines per Xe-core while Lunar Lake has eight 512-bit vector engines. While the total shader ALU count might be the same—you basically get 1024 shader operations per clock—Intel claims that Xe2 will deliver roughly 1.5X higher performance than Meteor Lake's Xe-LPG, thanks to improvements in various architectural enhancements.

Intel's next-generation Battlemage dedicated GPUs will aim to compete with the best mainstream graphics cards when they arrive in late 2024.

One of the other important changes is that Xe2 will also have eight 2048-bit XMX (Xe Matrix eXtensions) engines. The Arc Alchemist architecture includes XMX support on the dedicated desktop and laptop solutions, but it was removed from Meteor Lake's GPU to save on die space. Adding XMX back into the iGPU with Lunar Lake will mostly be useful for AI workloads, with up to 67 TOPS of compute, but things like XeSS upscaling can also benefit.

Hyper-Threading support (SMT), the first time a new desktop architecture has omitted that feature since it was reintroduced with Core i7 Bloomfield CPUs in 2008. Intel's rationale is that Windows often won't schedule tasks on the second thread for the P-cores until the E-cores have been filled, and the extra plumbing to make SMT work results in a less efficient design. Applications that can benefit from lots of threads will do better leveraging all the E-cores in terms of overall efficiency.

The number of P-cores and E-cores will remain the same as in the previous 14th Gen chips, topping out at eight P-cores and 16 E-cores. That means the total number of threads will drop from 32 on the 14900K to 24 on Arrow Lake, but presumably the architecture changes will more than make up the difference. Lesser models will have six P-cores, and some might have four P-cores, while E-core counts can be as low as four, increasing in multiples of four up to the maximum 16 E-cores.

The integrated graphics for Arrow Lake will continue to use the Xe architecture, though they'll get a bump up to 64 Xe-cores (compared to 32 in Raptor Lake). Clock speeds may also take a step backward from the 6.0 GHz Raptor Lake chips to a rumored 5.4 GHz peak clock on the fastest models.

Arrow Lake-compatible motherboards using socket LGA1800 were shown at Computex in June, with the chips expected to begin shipping this fall.

AMD THREAD SCHEDULING WOES

Some questions have come up with Zen 5 performance on certain CPUs. It goes back to AMD drivers that were created for the Ryzen 7950X3D and 7900X3D—the dual-CCD Zen 4 processors. Because only one of the CCDs has stacked cache, AMD's drivers work in concert with Windows 11's Game Mode feature to detect games and assign them to the appropriate CPU cores.

The driver uses a 'known good list' of games that benefit from the stacked L3 cache or higher clocks. If a detected game is on the list, the software will park the cores on the non-optimal chiplet, and target threads to the faster chiplet—that's usually the X3D CCD, but some games might prefer the higher clocks of the non-X3D chiplet. The same basic idea can also be used with the Ryzen 9950X and 9900X to improve communication latencies by keeping all threads for a game on the same chiplet, so AMD elected to use the same software on the dual-CCD Zen 5 CPUs.

It makes sense on one level, and it generally works as expected with a CPU like the 7950X3D or 9950X. There's one problem: once you've installed the necessary chipset drivers that enable the core parking and thread scheduling, you're stuck with them. The only way to get rid of them is to either fully reinstall the OS, or reimage the drive to an earlier state. This can create problems when the drivers reduce performance on single-CCD CPUs.

You'd think that AMD could simply disable the driver on CPUs that don't need it (ie. the Ryzen 5 and Ryzen 7 parts), but AMD says it's a Microsoft problem. Whatever the case, it seems like incredibly poor software design, something that should have been fixed soon after the feature became available. But here we sit, over 18 months later, and it remains a problem that may have impacted performance on some reviews of the 9600X and 9700X chips.

AMD 3D V-Cache Performance Optimizer Driver

Service
- amd3dvcacheSvc.exe creates the amd3dvcacheUser.exe process as user during session logon.

User Process
- amd3dvcacheUser.exe calculates a cpu ranking preferences based on EFFECTIVE_POWER_MODE_CALLBACK and EVENT_OBJECT_FOCUS events. On change, sends new preference to amd3dvcache.sys via DeviceIoControl.

Kernel Mode Driver
- amd3dvcache.sys communicates new preference to firmware via ACPI Device Specific Method.

Firmware
- AMD firmware Changes CPPC Highest Performance values and notifies OS of CPPC Highest Performance Change via a System Control Interrupt.

OS
- Thread Scheduling uses CPPC Highest Performance values which may change during runtime due to CPPC Highest Performance Change.

Originally introduced for 3D V-Cache Zen 4 CPUs, AMD's Performance Optimizer Driver and PPM Provisioning File Driver can cause reduced performance on CPUs that don't need them.

gamesradar+

One-stop destination for games, movies, TV, and more!

New Channels: Movies, TV and Cool Stuff!

www.gamesradar.com

Conclusion

YOU MIGHT THINK that after more than 40 years of x86 processors, there wouldn't be much room left for improvement, but AMD and Intel continue to find plenty of areas in which to innovate and improve with each passing generation. We'll see two new families of processors from both companies before the end of the year, with AMD's Ryzen 9000 and Ryzen AI 300 processors already available.

AMD's laptop chips look extremely promising, and while there's certainly a lot of hype behind the AI movement—which shows up here with a rebranding to Ryzen AI—there are major improvements on tap in the core CPU and graphics areas as well. It's interesting to note that neither Arrow Lake nor Granite Ridge, the desktop-centric CPUs from the two x86 juggernauts, are including NPUs. Presumably that's because power efficiency is less of a factor, and you can always add a GPU if you want fast AI performance. We'll have to wait and see whether Microsoft's Copilot features can make the leap from marketing hype to truly useful tools in the coming months.

The desktop Ryzen 9000 CPUs also look good, but AMD runs smack into a problem of its own making. It touted up to a 12 percent average improvement in gaming performance from Zen 5—when compared with the older Ryzen 7 5800X3D. Faster gaming performance would be nice, but benchmarks show that if you care about gaming, the Ryzen 7 7800X3D remains the fastest chip for now. We'll have to wait for Ryzen 9000 X3D CPUs to see how those stack up.

As for Intel, it has struggled of late, not just with CPU failures (see Tech Talk on page 11), but also with profits in general. Intel recently announced major cuts to its workforce, along with cuts to R&D—never a good thing for a company that strives to innovate. But what we really need to see is how Lunar Lake and Arrow Lake fare when compared with AMD's offerings, not to mention Qualcomm's Snapdragon X Elite SoC that's finally making Windows on Arm a good platform.

If you're not in a rush to upgrade or build a new PC, or to buy a new laptop, waiting until later this year to see how things shake down is the most sensible approach. Perhaps we'll even get some new GPUs to go along with our CPU upgrades before we say goodbye to 2024. All indications are that this year will end with a bang in terms of new PC hardware.

> We'll see two new families of processors from both companies before the end of the year, with AMD's Ryzen 9000 and Ryzen AI 300 processors already available

INTEL NPU 4.0

Lunar Lake will feature Intel's fourth-generation NPU, though we're not entirely clear on where the first two generations were used. It lists a 0.5 TOPS NPU 1 from 2018 as the first offering, which seems more like a CPU core than an NPU. There was a 7 TOPS NPU 2 in 2021 as well, but for most people, the first true NPU from Intel was the 11.5 TOPS model integrated into last year's Meteor Lake chips.

NPU 4.0 takes the base design of NPU 3.0, and then triples the number of compute units, boosting the maximum performance up to 48 TOPS—basically the same level of performance as AMD's XDNA 2 NPU in Strix Point. Obviously, there's more going on than just adding compute cores. Clock speed increases appear to account for another 40 percent boost, and Intel has reworked many other elements of its NPU design to improve performance and efficiency.

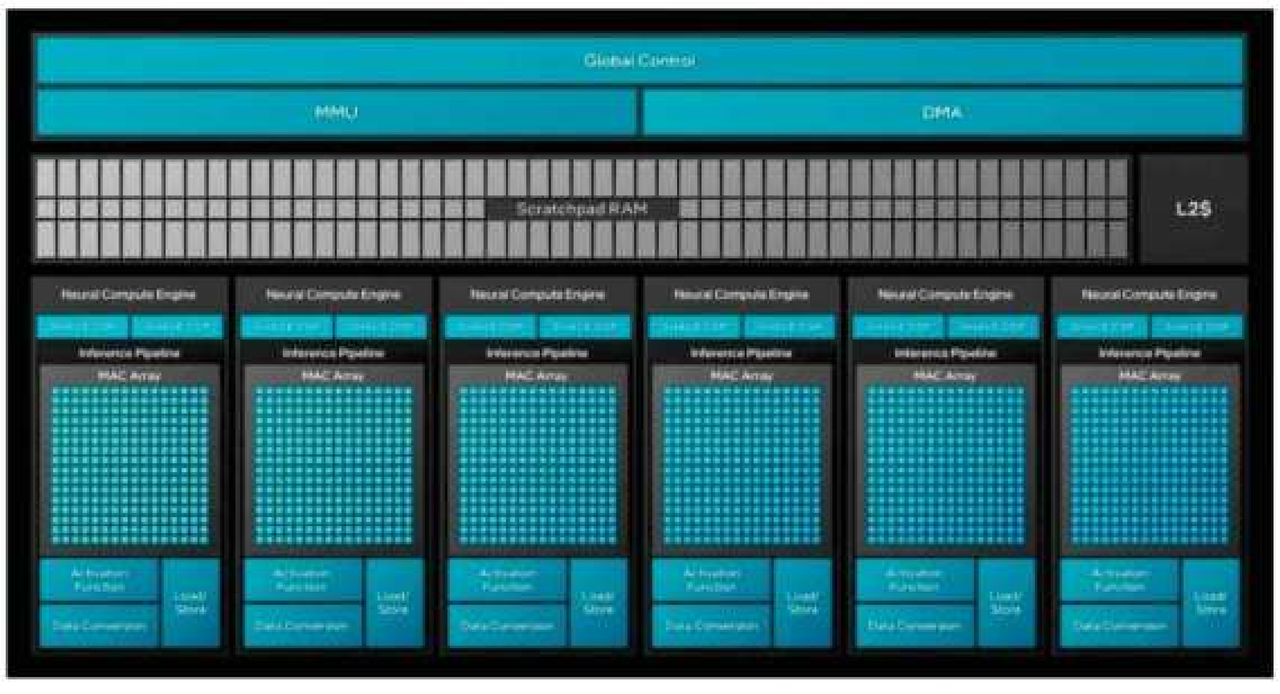

© INTEL

Efficiency is a key element of NPUs. Nvidia points out that if you're looking for maximum TOPS to run AI workloads, its GPUs have offered more performance than most NPUs for the past six years. The RTX 2080 offered up 170 TOPS of INT8 compute back in 2018—over three times the performance of the Strix Point and Lunar Lake NPUs. The difference is that NPUs will be able to deliver their 48–50 TOPS with a fraction of the power draw of a GPU—an RTX 2080 likely consumes around 120–150W to deliver its 170 TOPS.

Intel claims up to a 2X improvement in performance per watt for its NPU 4 compared to the previous NPU 3, with improvements in local bandwidth, the SHAVE DSPs, and Scratchpad memory. It has twice the memory bandwidth this generation, and as an example of the generational gains, Intel says that Meteor Lake's NPU could do 20 iterations of Stable Diffusion in 20.9 seconds using 9.0W of power on average, while the Lunar Lake NPU can complete the same task in just 5.8 seconds while averaging 11.2W, yielding a 2.9X improvement in overall efficiency.

Intel's latest NPU packs 12,288 MACs—Multiply and Accumulate units—triple the number found in Meteor Lake, for triple the performance.

SUBSCRIBE TODAY

Find your next great read

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

www.magazinesdirect.com/B3XMPC

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

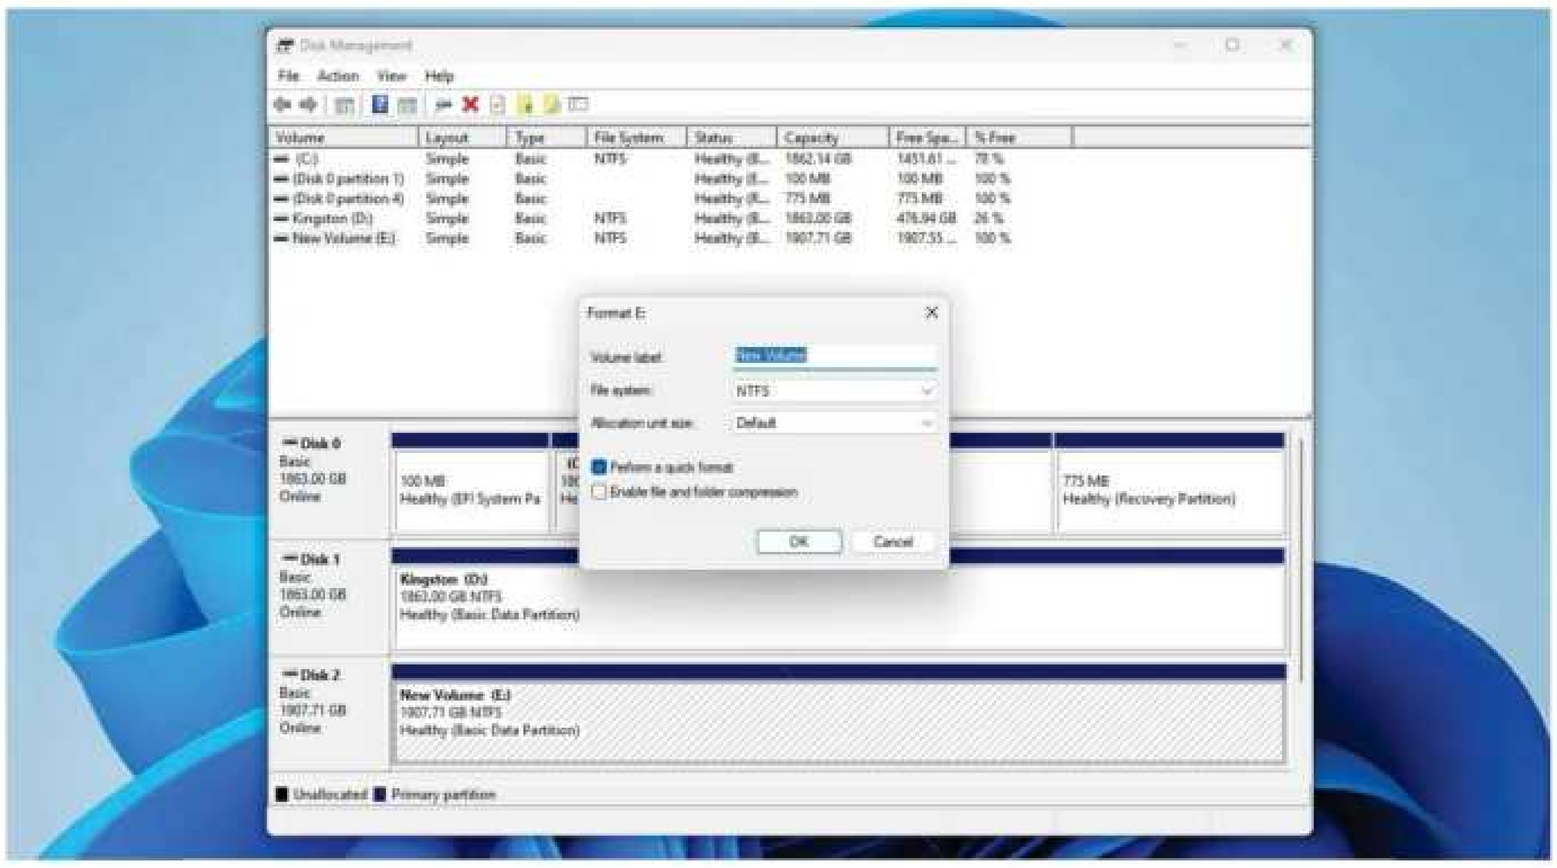

ALWAYS PERFORM A QUICK FORMAT

Whatever you do with your SSDs, ensure that when formatting, you perform a quick format. If you try to do a full format on an SSD, it treats it like a traditional HDD, and Windows will run a full read/write cycle over every cell on the SSD, writing a 0 in every memory cell before deleting it. If your SSD only has a 600 TBW endurance rating, and it's a 2TB drive, then you've immediately reduced that to 598 before writing anything to it.

ZAK STOREY
CONTRIBUTOR

FROM SCEPTIC TO CONVERT

I've always been a bit fascinated by *Blade Runner 2049*. One of the central characters is a holographic AI by the name of Joi. She acts as a romantic companion to the protagonist, K, throughout the film, as a positive, encouraging presence.

Now, that clearly poses some ethical quandries. Joi is effectively the property of K, who in turn is property of the police department, but it got me thinking more in line with our own expectations of AI.

If you speak to any character AI, or LLM-based ANI, you'll know that they often have follow-up queries. In a lot of ways, it is possible for them to behave in a similar manner to Joi. I did some research, and found a number of sites that operate AI partners.

Again, there are some ethical questions, but it gives me hope. One of the biggest killers out there is loneliness. In fact, there have been multiple studies indicating that loneliness is a powerful indicator of an increase in all-cause mortality.

I'm not suggesting that AI should or could replace human interaction, but if it could alleviate some of those symptoms and provide some comfort, well, it's one of the things that swung me from AI sceptic to convert. Yet, the fact that it's still not legislated does concern me.

MAKE – USE – CREATE

62
Lock down your browser

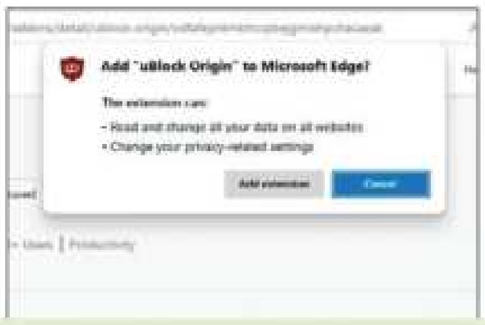

66
Clear your PC of stalkerware

68
Boost your gaming VPN

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

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

HMD Skyline

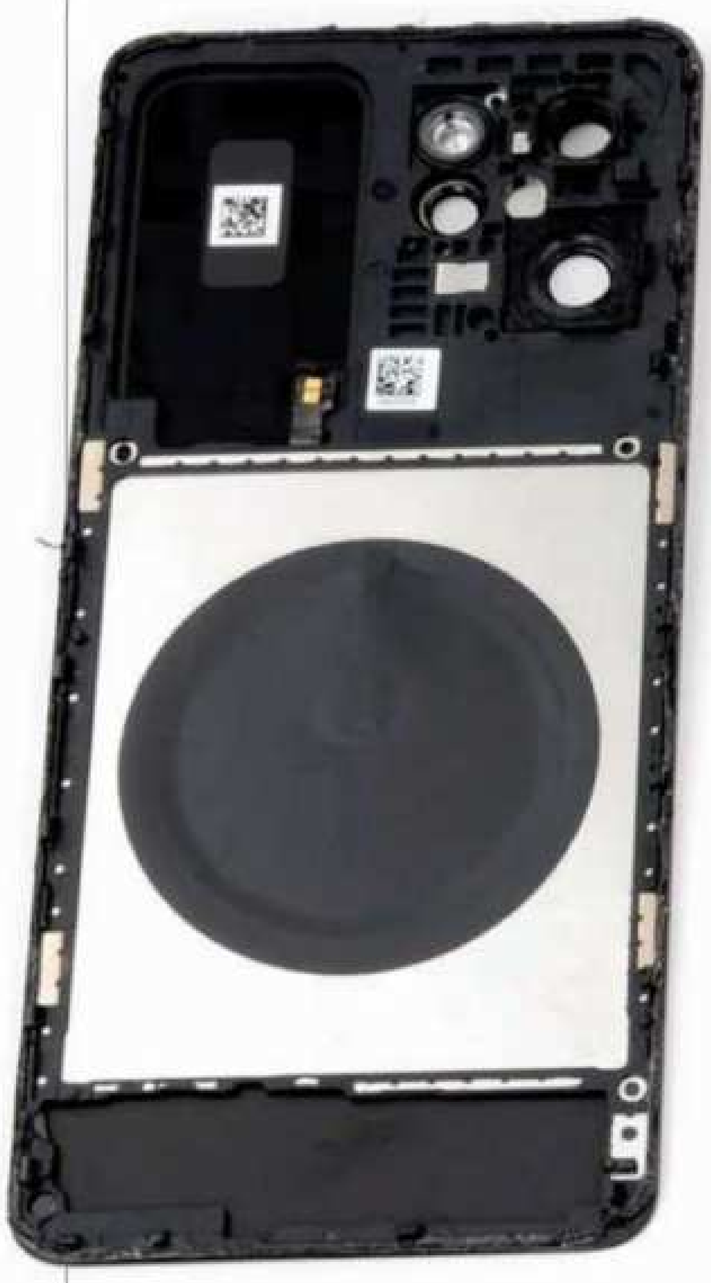

It's not the easiest replacement job ever, but at least with some patience and solvent, you can fit a new battery yourself.

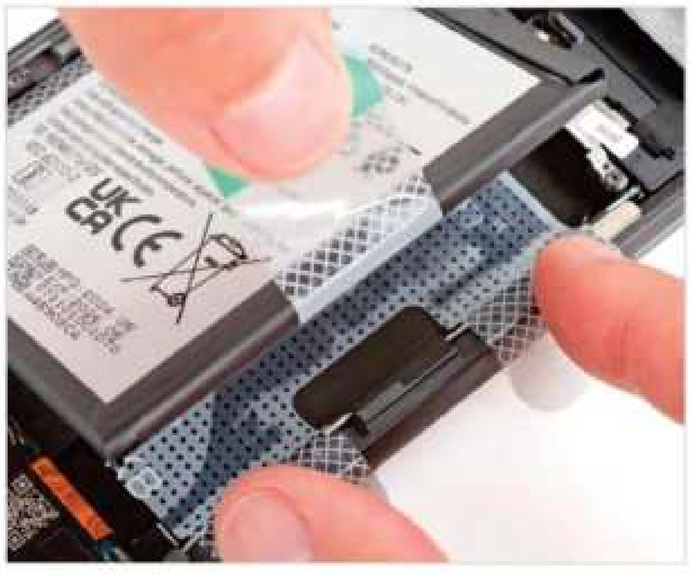

Simply turn one screw to pop open the back cover of the HMD Skyline.

About iFixit

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

If you want to surprise and delight a phone repair tech or teardown engineer, this is how you do it.

Remove a few screws, and the screen drops out freely from the other side. Rarely have screen repairs been this easy.

BACKGROUND

Finnish mobile maker HMD, which owns and manufactures Nokia handsets, has made a pledge to offer more repairable phones. The Skyline aims to deliver on this promise.

MAJOR TECH SPECS

- Screen: 6.55 in 20:9, OLED, 144Hz refresh rate, 1000 nits peak brightness, FHD+ 1,080 x 2,400
- Cameras: Front: 50 MP, Rear: 108 MP Main AF with OIS, 13 MP, 50 MP Tele
- Connectivity: Bluetooth 5.2 with LE Audio, USB Type-C, Analog audio output, WiFi 6e, NFC, 5G
- Battery: 4,600 mAh (replaceable), 33W charging support (QC4.0 and PD3.0 PPS compatible), 15W magnetic wireless charging, 5w reverse wireless charging, Qi2 certified
- Memory/storage: Internal: 128/256 GB, MicroSD card support up to: 512 GB, RAM: 8/12 GB RAM
- CPU: Snapdragon 7s Gen 2
- Operating System: Android 14
- Security: Fingerprint sensor (Side Power key), Face Unlock, three years of security updates
- Materials: Tempered glass back, 100 percent recycled aluminum frame, IP54 IP rating

KEY FINDINGS

- HMD made headlines in 2023 with the Nokia G22. "To make people want to repair and keep their phones, we have to make the experience of repairing more compelling than switching to a new one," the company said. "So for our future devices, we are looking to both improve the ease of repairability, but also the experience of repairing your phone."
- Designed from the ground up to prioritize user-friendly entry and no-nonsense servicing, it shows that HMD is taking its repairability commitment seriously, and having fun doing it.
- There's nothing in the industrial design to give away the fact that there's repairable hardware underneath. On the bottom edge lies a Torx screw—unassuming, chassis-colored, and well-integrated into the design aesthetic. But with a counter-clockwise twist, it pops up the corner of the back cover—just enough for you to slip a pry tool underneath.
- The back cover uses adhesive for watertightness, but it yields to prying and slicing. When you're done, it clips back together. We still wish for a modern waterproof smartphone with reusable gaskets instead of adhesive—but for now, this is as clean a design as we've seen.
- The battery remains adhered in place, but comes with a built-in plastic liner that you can pull apart to separate the battery from the chassis. Before installing a new one, you have to clean out the adhesive residue, and install fresh adhesive strips. This requires patience and some solvent. We prefer batteries secured with reusable fasteners or hard-shells that you can pop in and out in seconds, but if you must rely on adhesive, this is among the better designs.
- The remainder of the Skyline's internals are fairly straightforward to navigate, with a high degree of modularity. The daughterboard, underneath a modular loudspeaker, houses the USB-C port and microSD/SIM card combo tray. The cameras and mechanical button repairs aren't well optimized, and require a lot of disassembly, but are on par with similar phones in this category.
- Repairability Score: 9 out of 10 (10 is easiest to repair). A clever entry, a repair-friendly screen, straightforward battery replacement, and public availability of parts all help land the Skyline near the top of the leaderboard.

Lock down your browser

YOU'LL NEED THIS

A STABLE INTERNET CONNECTION

ADAPTIVE AND EVASIVE MALWARE, phishing domains, harmful links, and packet sniffing can all target a common point of weakness: your web browser.

In May, Google was forced to release an emergency security update for the Chrome browser, which exploited a bug in the V8 JavaScript engine to allow remote code execution attacks. In August, vulnerabilities were also discovered in Edge that targeted out-of-bounds memory access in the ANGLE graphics backend to divert users to malware domains.

2024 has also seen a spike in 'malvertising' whereby legitimate-looking web page ads lead to malware. Given that browser developers are constantly releasing patches, it's obvious that regular updates are one of the best ways to surf safely.

Still, by default, most browsers display online ads without first checking their legitimacy. Sites are free to download tracking and session cookies to monitor your online activity. Many programs also send User-Agent request headers to sites to make it easy to 'fingerprint' your browser and device. Although connections to many sites are secured by TLS, bad actors can easily monitor which sites you've visited through your browser's unencrypted DNS requests.

In this guide, we'll explore how to choose a secure browser. You'll also learn how to configure it to maximize your online security while web surfing. **–NATE DRAKE**

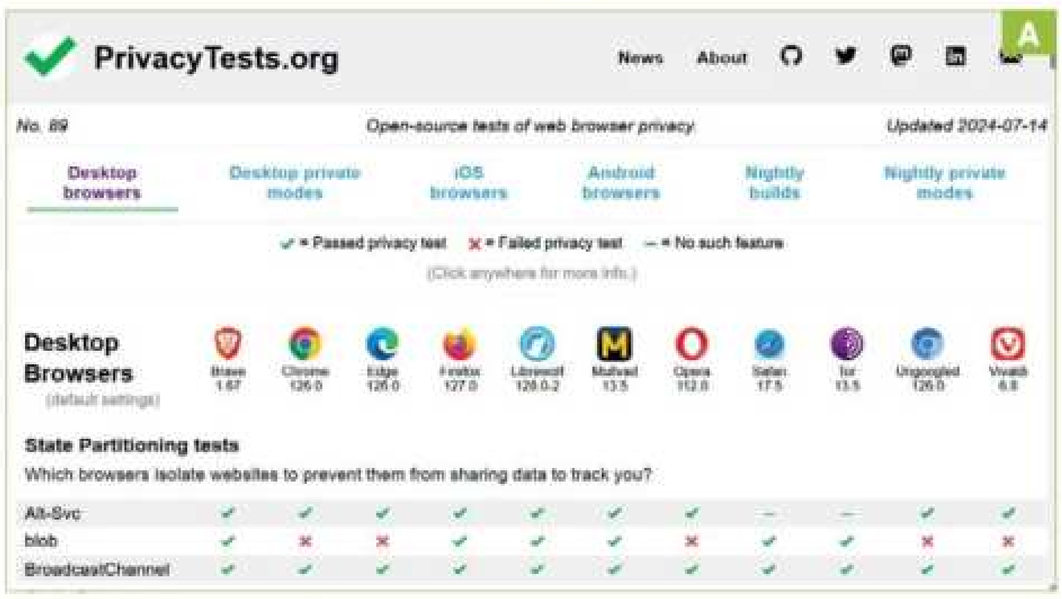

1 BEST OF BROWSERS

At just under 66 percent, Chrome dominates the browser market, but this doesn't make it the safest browser [Image A]. In late 2023, Google settled a multi-billion dollar lawsuit over alleged user tracking, even when the browser was in Incognito mode. The Manifest V3 API will also make it harder for ad blockers to function on Chromium-based browsers.

» The most secure browsers are fully open source. In the words of Linus Torvalds, "Many eyes make bugs shallow." In other words, security bugs are easier for the community to detect. If an open-source browser introduces an unpopular feature like Manifest V3, developers can create their own 'fork'.

» The Chromium web browser, on which Chrome is based, is free and open-source. Many proprietary browsers like Chrome, Edge, and Opera use it as a code base. However, there are also open-source Chromium-based browsers that include more secure features.

» For example, both Iridium and Brave have removed any code that 'dials home' to Google. Both browsers also incorporate features to block ads and third-party trackers.

» While Mozilla Firefox doesn't enjoy the popularity it used to, it also has some excellent features for safer web browsing, including enhanced tracking protection, DoH (DNS over HTTPS), and fingerprinting protection.

» The Firefox fork LibreWolf includes all of the above, with extra security enhancements, such as disabling cloud sync and sponsored shortcuts.

» We've not recommended a specific open-source browser, as it's important to research one that suits your needs. If you need help, visit https://privacytests.org for a comparison of privacy features in major browsers.

» For the purposes of this tutorial, we will focus on locking down a standard install of Firefox.

2 CONFIGURE BROWSER PRIVACY FEATURES

Despite Firefox's impressive array of security features, by default, the browser can run certain tests and collect data on your computer.

» To fix this, go to Settings > 'Privacy & Security'. Next, and scroll down to 'Firefox Data Collection and Use'. Here, you can untick the boxes to allow sending data back to Mozilla, as well as to install and run 'studies'.

» Once this is done, scroll to the 'Browser Privacy' setting at the top. 'Standard' protection is usually sufficient for blocking most trackers, cross site cookies, cryptominers, and fingerprinting code. You can, however, enable 'Strict' mode to disable suspected fingerprints at the risk of causing some sites to load incorrectly.

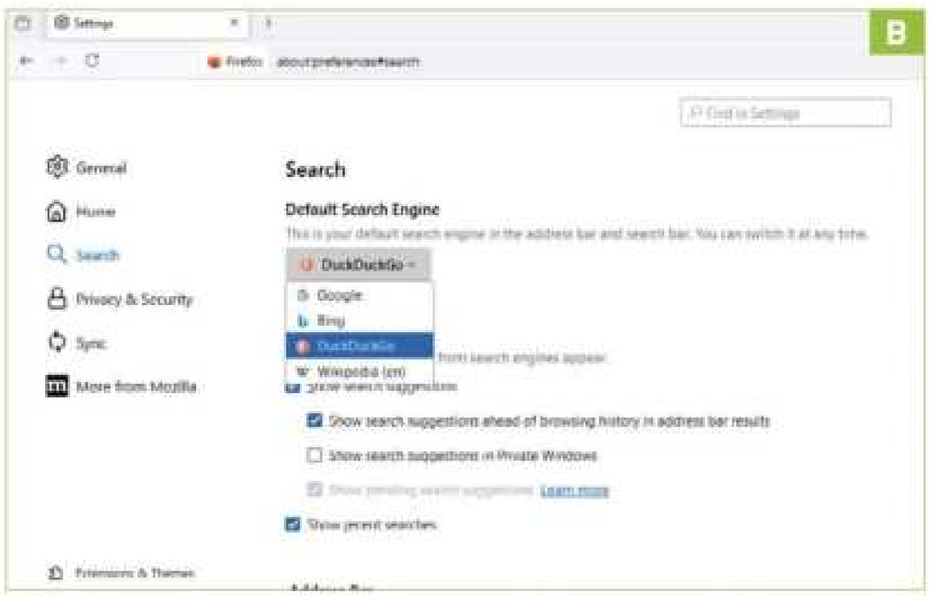

©PRIVACYTESTS.ORG; MOZILLA

A SECURE BROWSER

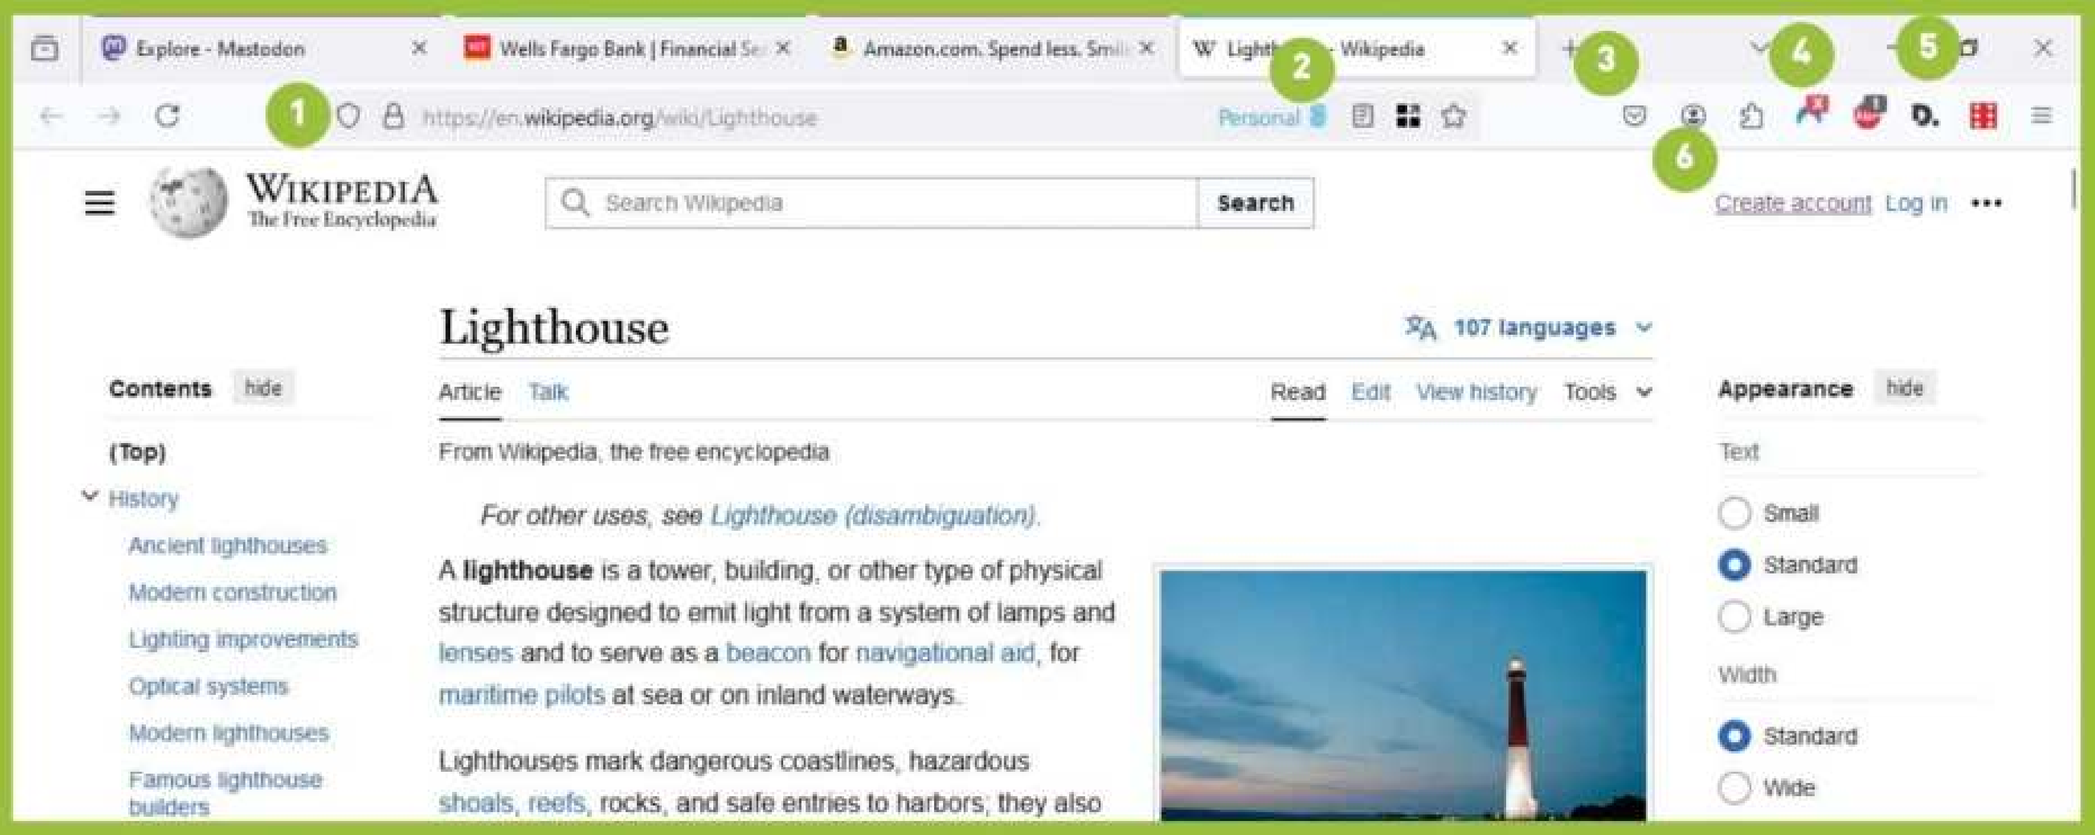

1. ANTI-TRACKING FEATURES
The best browsers integrate features to automatically detect and prevent social media trackers, cross-site cookies, and cryptominers. You can also make exceptions so that trusted domains load correctly.

2. CONTAINERIZED BROWSING
Firefox containers isolate browsing sessions using color-coded tabs. Web data in one container cannot be accessed by others. You can even sign in to multiple accounts on the same site.

3. FINGERPRINT RESISTANCE
Some browsers, like Brave, have built-in features to provide websites with randomized fingerprints. For others, you can disable features that could be exploited by fingerprinting scripts.

4. VPN
While we recommend using client software over a browser extension, VPNs provide a secure connection between your device and the VPN server. As your traffic is encrypted, it's almost impossible to monitor.

5. SECURE EXTENSIONS
Ad Blockers like AdBlock Plus can block targeted ads from loading in web pages. Other extensions, like Disconnect, can prevent trackers. Add-ons need to be configured to be fully effective.

6. LOCAL PASSWORD MANAGEMENT
Many browsers support syncing your passwords across devices, but it's safer to store credentials locally. Ideally, these should be encrypted with a master password. Use Diceware for password generation.

» Like most browsers, Firefox also has the option to ask websites to include a 'Do Not Track' request in HTTP headers. Remember, there's no way for your browser to enforce this rule.

» From here you can also have Firefox delete cookies and site data each time the browser closes, though it's far better to use a browser's 'Private' or 'Incognito' mode, which we'll explore later.

» The default search engine for most web browsers is Google. While it has an excellent reputation for serving relevant content, alternative search engines like DuckDuckGo don't track your search data to serve you targeted ads [Image B]. DDG also blocks tracking code in embedded social media posts.

» To make the switch in Firefox, select the 'Search' section from the left-hand pane, then choose an alternative from the drop-down menu under 'Search Engine'. From here, you can also disable 'Show Search Suggestions' based on your browsing history.

» These days, browsers are much better at enforcing HTTPS. In other words, the software will access only secure versions of websites protected by SSL/TLS. When there's both a secure and a non-secure version of a website, hackers sometimes try to direct your device to the latter in order to monitor your traffic.

» To prevent this in Firefox, return to 'Privacy & Security', and scroll down to 'HTTPS Only Mode'. Choose to enable it for all windows.

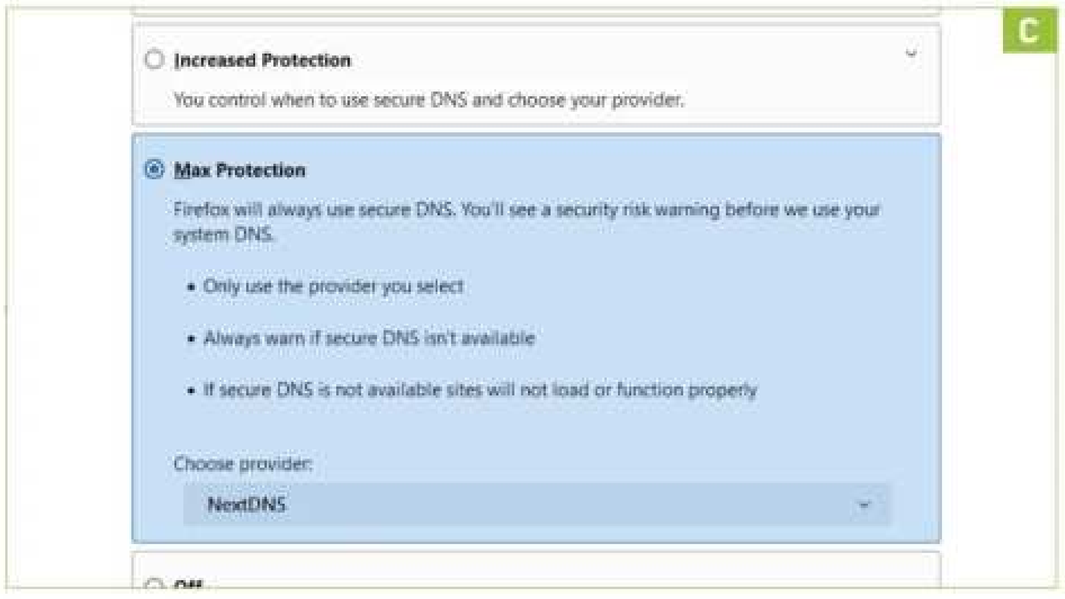

3 CONFIGURE DNS AND HISTORY SETTINGS

The DNS (Domain Name System) acts as a virtual phonebook to convert human-readable domain names into machine-readable IP addresses. Most internet users have a local 'resolver' provided by their ISP to handle DNS queries. [Image C]

» By default, these are unencrypted, so it's easy for bad actors to monitor the websites you visit. Some cybercriminals also engage in 'DNS Poisoning', whereby they set up their own DNS resolver to redirect your browser to malicious 'phishing' domains designed to resemble legitimate sites.

» DNS over HTTPS (DoH) offers an excellent solution by encrypting your DNS queries in the same way as for

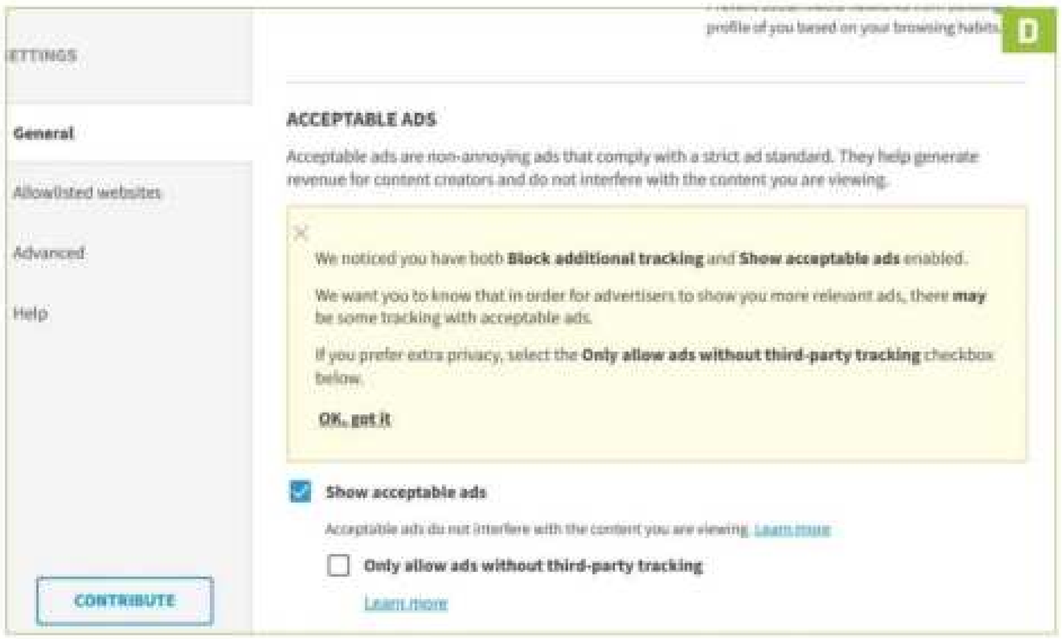

a web page secured by an SSL certificate. Its implementation varies from browser to browser.

» For instance, Firefox's 'Privacy and Security' section uses 'Default Protection'. In plain English, this means the browser will use DoH in supported regions, but default to your regular DNS resolver if it's unable to do so.

» To ensure you're always using Secure DNS, scroll down to 'Enable DNS over HTTPS using:' and select 'Max Protection'. This will require Firefox to always use DoH. Use the drop-down menu to select from one of the preconfigured providers (Cloudflare or NextDNS), or choose 'Custom' to enter your own.

» Even if your DNS queries can't be detected by your ISP, by default, your browser will store visited websites in its cache, along with cookies and other temporary files. This means anyone with access to your device can monitor your browsing activity.

» If you use a 'private' or 'incognito' window to access sensitive websites, your browsing history won't be saved, and any temporary files will be cleared when you close the browser.

» You can also configure your browser never to store your browsing history. In Firefox, there's a check box in the 'History' section of 'Privacy & Security' labeled 'Always use Private Browsing Mode'. Restart the browser to enable this.

4 ENABLING EXTENSIONS

You should only have one browser extension for each purpose. For instance, you're unlikely to gain better protection by choosing two ad-blockers simultaneously.

» Your extensions may need some tweaking for you to benefit from them fully. For example, AdBlock Plus (https://adblockplus.org) allows some 'non-intrusive' advertising'.

» To fix this, you need to access extension settings once AdBlock Plus is installed, then uncheck 'Show Acceptable Ads' [Image D]. Alternatively, just install uBlock Origin (https://ublockorigin.com), which blocks virtually all advertising out of the box.

» For the same reason, we recommend installing Disconnect (https://disconnect.me) over Ghostery. While the latter is better known for blocking trackers, Disconnect doesn't send telemetry to the developers.

» If you followed our suggestion to use Firefox, you can also take advantage of container tabs. This allows you to separate your browsing experience into groups of colored tabs. Each container is isolated, so any browsing data saved in one container session can't be accessed by another. This means you can sign in to multiple accounts on the same website using different container groups.

» To get started, install the Multi-Account Containers extension (https://mzl.la/3SNjid3). Next, go to Firefox settings and find the 'Tabs' section. Click 'Settings' to view the default container tabs.

» Click 'Add New Container' and enter a relevant name, eg. 'social media'. Choose a fitting color and icon, then hit 'Done'.

» Open a new tab and go to an appropriate website for your new container, like https://mastodon.social/explore. Next, right-click the tab itself and choose 'Open In New Container Tab' > 'Social Media'. The website will open in a new tab color-coded for that particular container. The container name will also be displayed in the address bar. Take some time to open tabs for other categories like 'Shopping' and 'Banking' to get used to working with containers.

5 ENCRYPTING YOUR CONNECTION

While SSL/TLS can do wonders to encrypt traffic between your device and websites, it's still trivial for anyone with access to your ISP records to know what sites you've visited.

THE MANIFEST V3 CONTROVERSY

In late 2021, Google announced plans to deprecate its Manifest V2 API. As of June, browser extensions using Manifest V2 are no longer supported.

This change has proved controversial, as Manifest APIs manage how extensions interact with Chromium-based browsers like Chrome, Edge, and Opera. V3 has introduced some changes, supposedly to make browsers safer and improve user experience. This means that all browser extensions must contain all the code they run. Extensions will require permission from Google to implement any in-browser changes.

The consequences for ad blockers could be disastrous. Such extensions work by downloading lists of HTTP requests. Accessing and updating these quickly is an essential part of keeping ad blockers functioning.

Manifest V3 forces browser extensions to use declarativeNetRequest API, with a limit of 30,000 coded 'rules'. As most ad blockers need at least ten times this, they won't run well in Chromium-based browsers. Google relies on targeted advertising, so it's easy to be cynical about Manifest V3.

These changes won't affect ad blockers in non-Chromium browsers like Firefox, as Mozilla currently has no plans to deprecate Mozilla V2. Some Chromium-based browsers, like Brave, have built-in ad blocking features, so also won't be affected.

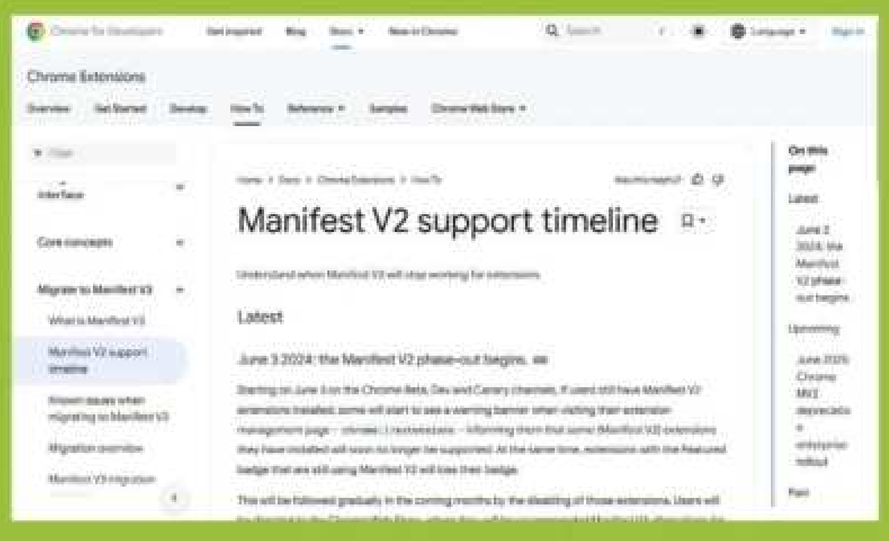

© GOOGLE, ADBLOCK PLUS

» The Tor Browser (www.torproject.org) is a fork of Firefox designed to route traffic through the 'darknet'. As your data is encrypted and routed through multiple 'relays', it's almost impossible for others to trace your device location or read your internet traffic.

» Unfortunately, routing traffic through multiple Tor relays can significantly slow page loading times. If you're accessing regular 'clearnet' websites, your traffic also can be intercepted and read by owners of 'exit relays'. You can sidestep this issue by only accessing Tor hidden services (.onion addresses).

» Using a VPN (Virtual Private Network) offers much better surfing speeds with comparable anonymity. If correctly configured, your device can establish an encrypted connection to the VPN server [Image E]. This means that anyone with access to your ISP's records won't be able to view your browsing history. Most modern VPN services also manage DNS requests. You can double check this is happening correctly by visiting https://ipleak.net.

» If you plan to use a VPN with your browser, we recommend installing dedicated client software rather than a browser extension, as this way, all your internet traffic is protected. If you do install the extension, make sure you enable it for 'Private' windows. In Firefox, you can also right-click the extension to 'Pin to Toolbar'.

» Recommending a specific VPN service is beyond the scope of this tutorial, but our sister site, TechRadar, maintains a list of some of the most popular: (www.techradar.com/vpn/best-vpn).

» The best VPN services regularly submit to an audit of their 'no log' policy, to independently verify that they don't monitor your internet activity.

6 RESIST BROWSER FINGERPRINTING

When connecting to a web server, every modern browser sends a User-Agent request header. This header includes information like your browser version and screen resolution. Sites can also inspect installed plugins and other types of software configuration to build a 'unique' fingerprint of your browser to identify you. As this information is stored server-side, using 'Incognito mode' or changing your IP address won't protect you from being profiled in this way.

» The Tor Browser uses several techniques to resist fingerprinting such as 'letterboxing'. This technique adds margins to the browser window to make it harder to detect screen size.

» The browser also uses NoScript (https://noscript.net). This extension, which is available for all major browsers, blocks scripts running on web pages, allowing users only to manually enable legitimate ones. As most types of fingerprinting rely on JavaScript, this can stop it in its tracks, but NoScript can prevent pages from loading properly unless you enable the right code.

© NORDVPN/IPLEAK.NET, MOZILLA

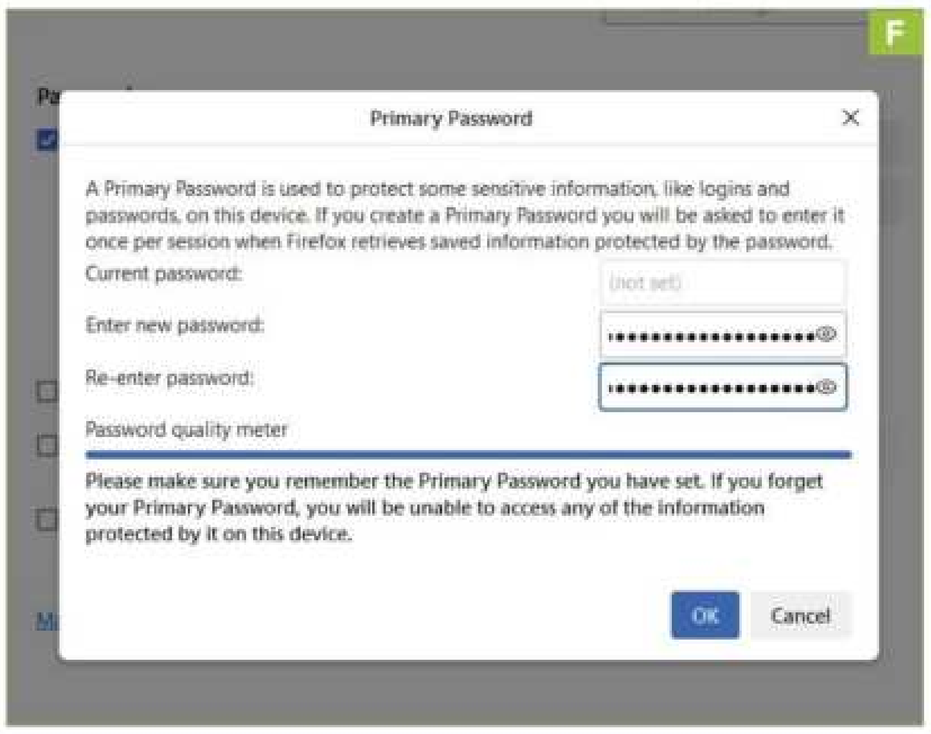

» Canvas Fingerprinting is a very common method to identify your device. It uses the HTML5 canvas element present, which is deployed by the Web GL JavasScript API to render 2D and 3D graphics in your browser. The fingerprinting technique involves a site inserting invisible text or images, then formulating a hash of how your browser renders the pixel data. As this will be slightly different, depending on your software, GPU, and graphics drivers, it's easy to identify specific devices.

» If you don't want your browser to be tracked across websites and sessions in this way, the easiest fix is to disable WebGL from your browser. In Firefox, you can do this by typing 'about:config: into the address bar. Next, search for 'webl.disabled', and change 'false' to 'true'.

» If you use Brave, Canvas and WebGL are already blocked for third-party APIs. The browser also uses a technique called 'farbling' to provide randomized values for common APIs, making it harder to fingerprint consistently.

7 PASSWORD PROTECTION

As impressive as online password managers are, there's a risk in trusting your credentials to them. LastPass users discovered this the hard way in late 2022 when they found certain customer data had been stolen. Since many password management platforms are closed-source, it's also difficult to verify that credentials are being encrypted properly.

» Using an open-source browser password manager allows you to store your credentials locally. Of course, this isn't much good if a bad actor accesses your device.

» Firefox resolves this by letting you set a 'Primary' password. This is required the first time in each web session to access the credentials database. To set this up, go to Firefox settings > 'Passwords'. Click the ... at the top right, then 'Options'. Check the box marked 'Use a Primary Password', [Image F], then enter a passphrase.

» Most browsers allow you to sync passwords across devices. Firefox handles this by creating an online account, into which you can sign in. (Remember though you'll need to set a new primary password.) Brave automatically generates a 'sync chain' series of words that you can enter into other browser installs to load your credentials.

» Most browsers can also suggest unique passwords, though they aren't always very complex. We recommend using Diceware (https://diceware.dmuth.org) for generating high-entropy passphrases.

Remove stalkerware from your PC

YOU'LL NEED THIS

A WINDOWS PC

ACCORDING TO KASPERSKY'S LATEST 'State of Stalkerware' report, over 40 percent of those surveyed worldwide said they'd experienced stalking or suspected that they were being stalked.

Stalkerware is a form of spyware, which is used by bad actors like domestic abusers to monitor people's devices without their knowledge. While it's most commonly found on mobile devices to harvest information like GPS data, your PC can also be infected.

Keyloggers, which silently record your keyboard activity, are a good example. Other forms of stalkerware may monitor your personal files, or even spy on you via your webcam.

Stalkerware is difficult to remove, as it's deliberately designed to disguise itself as a legitimate program or system process.

In this guide, you'll discover how to reboot your machine to safe mode, so the stalkerware can't execute. You'll also discover how to prevent it from launching, as well as how to remove malicious files from your web browser and system cache. **–NATE DRAKE**

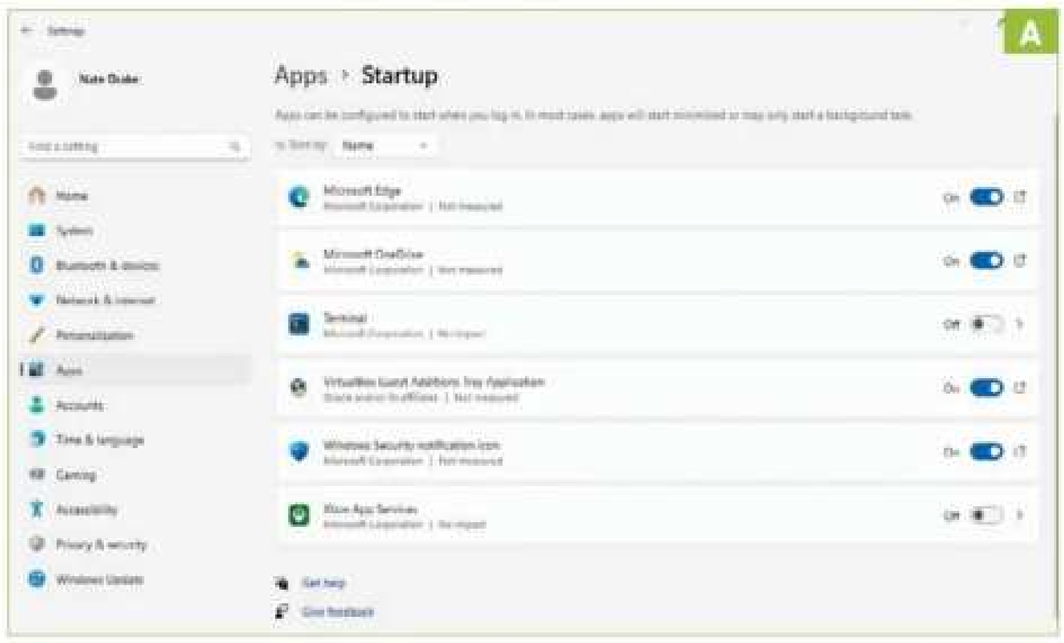

1 PLAYING IT SAFE

If you suspect that your machine has been infected, restart it immediately. On the sign-in screen, hold Shift to restart again. Next, choose 'Troubleshoot' > 'Advanced Options' > 'Startup Settings' > 'Restart'. Tap '4' to boot into safe mode.

» This mode starts with a minimal set of applications and drivers. This means that stalkerware can't interfere with your attempts to remove it.

» It's unlikely that malware will hide anywhere obvious, but to be on the safe side, type 'Installed apps' into the search bar to view the relevant section in Windows settings. Click the ... options icon, and uninstall any you don't recognize.

» Next, enter 'startup' into the search bar to view 'Startup Apps'. Click the rocker switch to deactivate any unwanted programs. [Image A]

» Malware can also use registry entries to launch undetected on startup. Enter 'regedit' into the search bar to open the Registry Editor.

» Navigate to HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion. Click 'Run' to view apps programmed to launch on system startup.If any are unfamiliar, right-click the registry key name, and choose 'Delete'.

» While this can prevent installed malware from loading, malicious files can also lurk in your browser's temporary files, cookies, and even in extensions. Open File Explorer, and enter the following address:

```
C:\Users\(yourusername)\AppData\Local
```

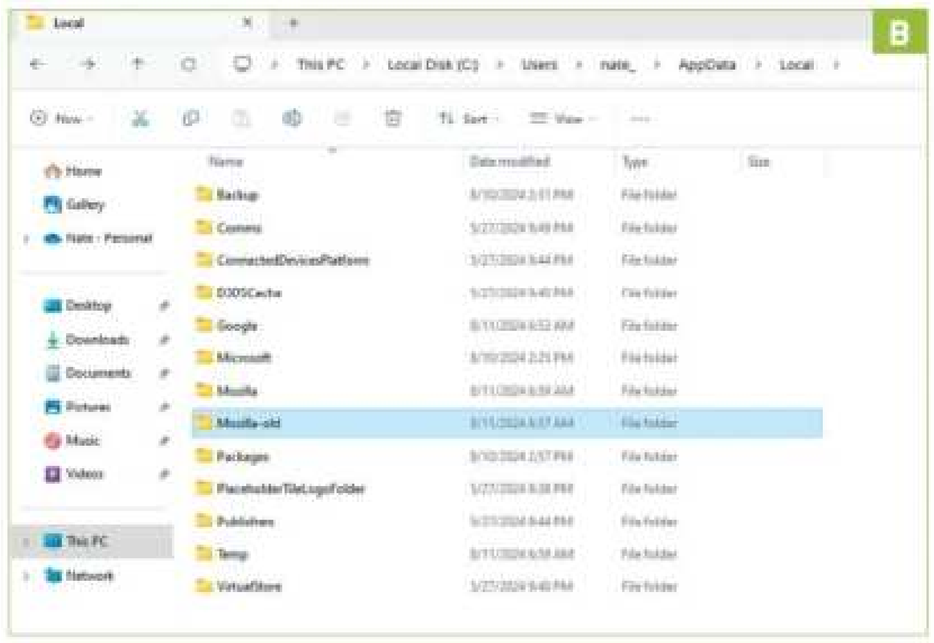

» This folder contains directories with your specific browser settings. In the case of Edge, these settings can be found in 'Microsoft' > 'Edge'. Chrome settings can be found in 'Google' > 'Chrome', Firefox settings can be found in 'Mozilla' > 'Firefox', and so on.

» Technically, you can delete these folders to reset the browser to its default settings. However, if you're concerned about losing important data, you can rename the directory, eg. from 'Mozilla' to 'Mozilla-old'. Next time you launch the browser, it will reset to its default settings and create a new folder. Your previous browser data will be preserved in the 'old' folder. [Image B]

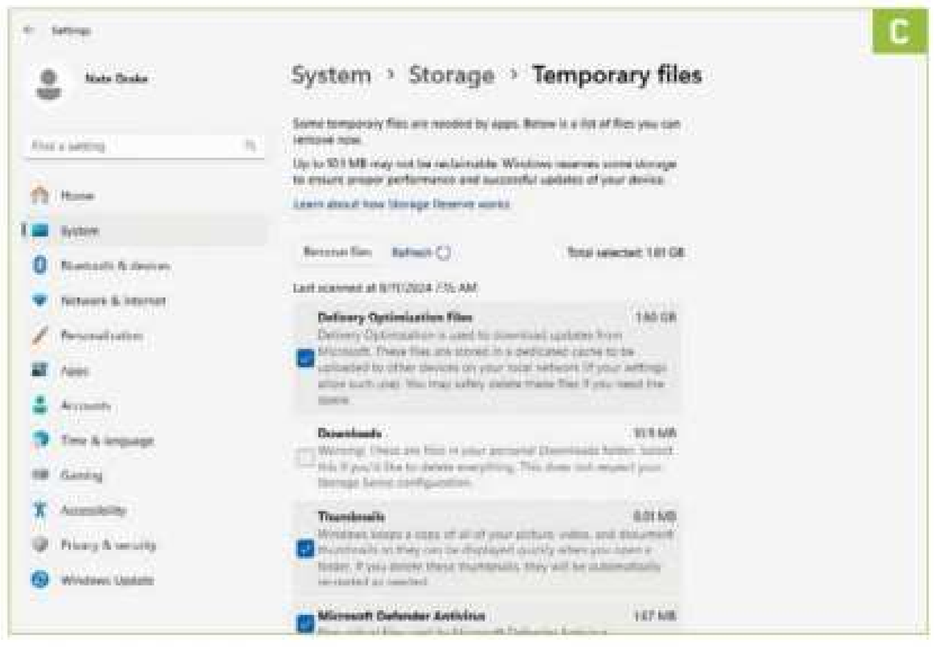

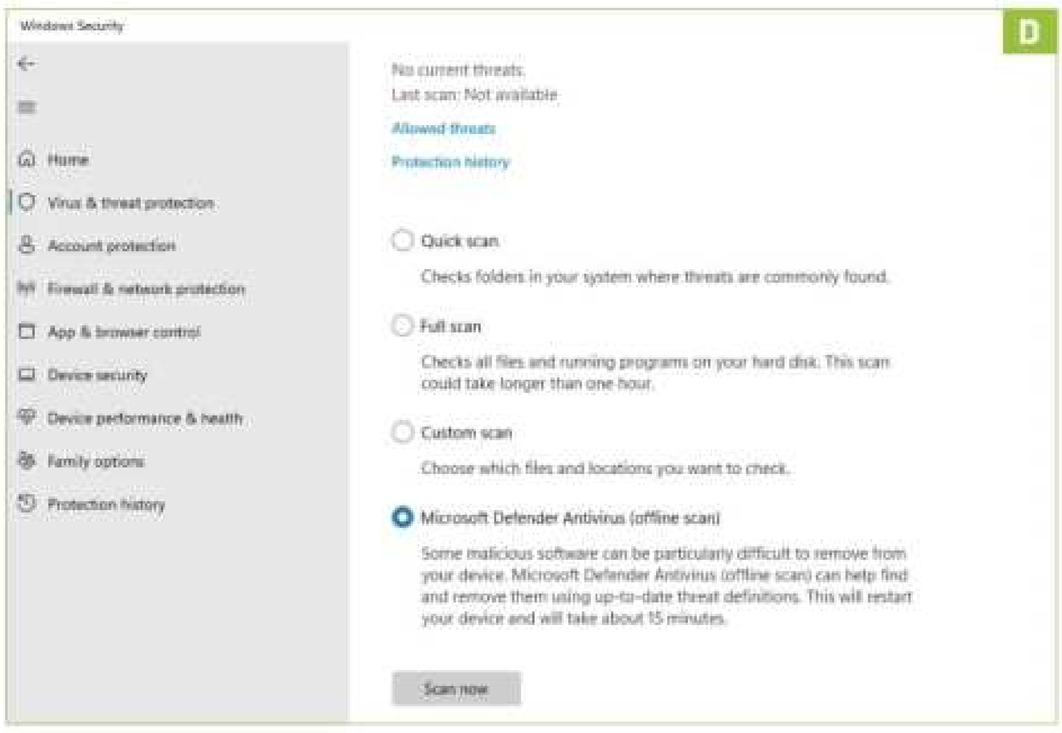

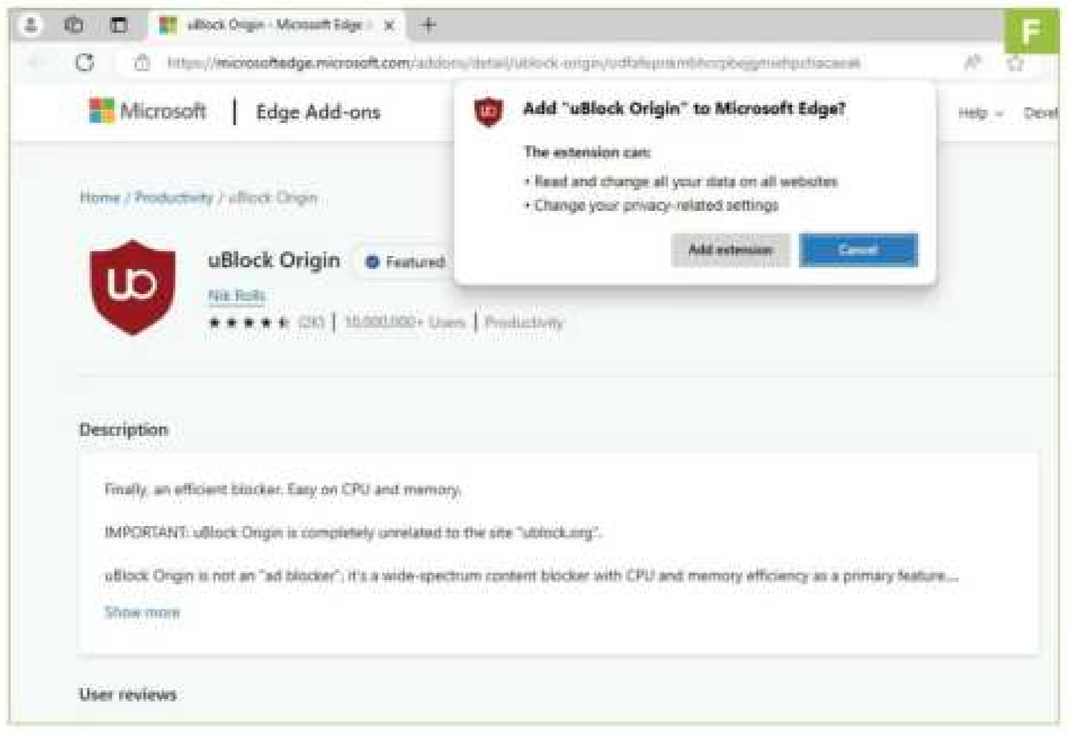

2 CLEANING UP

Stalkerware can still lurk in temporary files. To clear these, you'll need to reboot Windows back into standard mode. On login, enter 'storage' into the Windows search bar to view 'Storage Settings'. Next, click 'Temporary Files'.

» As you'll see, by default, both temporary internet files and temporary files generated by apps are covered in the cleanup. You can also opt to remove files from your personal 'Downloads' folder and 'Recycle Bin', if you wish. Make your choice, then select 'Remove Files' to continue. [Image C]

» If you have any files that aren't backed up, now is a good time to connect an external drive and copy them over. Remember, only do this for non-executable files like documents and images.

» Next, enter and select 'Windows Security' in the search bar. Click 'Virus & Threat Protection', then 'Scan options'. From here, choose 'Microsoft Defender Antivirus (Offline Scan)'.

» As the description explains, this is an excellent way to remove persistent malware as it updates virus definitions, then restarts the device so stalkerware won't be active during the scan. Click 'Scan now' > 'Scan' to proceed. [Image D]

» Once your system reboots, enter 'uninstall' into the Windows search bar to launch 'Add or Remove Programs'. It's unlikely that stalkerware would be displayed here, but play it safe by inspecting each of these applications and clicking ... (options) > 'Uninstall' next to any you don't recognize.

» Restart your machine, then use Ctrl + Alt + Del to launch Task Manager.

» Click the 'Startup Apps' icon in the left-hand pane to check once again that no unrecognized applications are listed. Next, return to processes, and check each one manually to be sure that you recognize it.

» If you see anything suspicious, right-click to 'open file location' and see where the program files are based. [Image E]

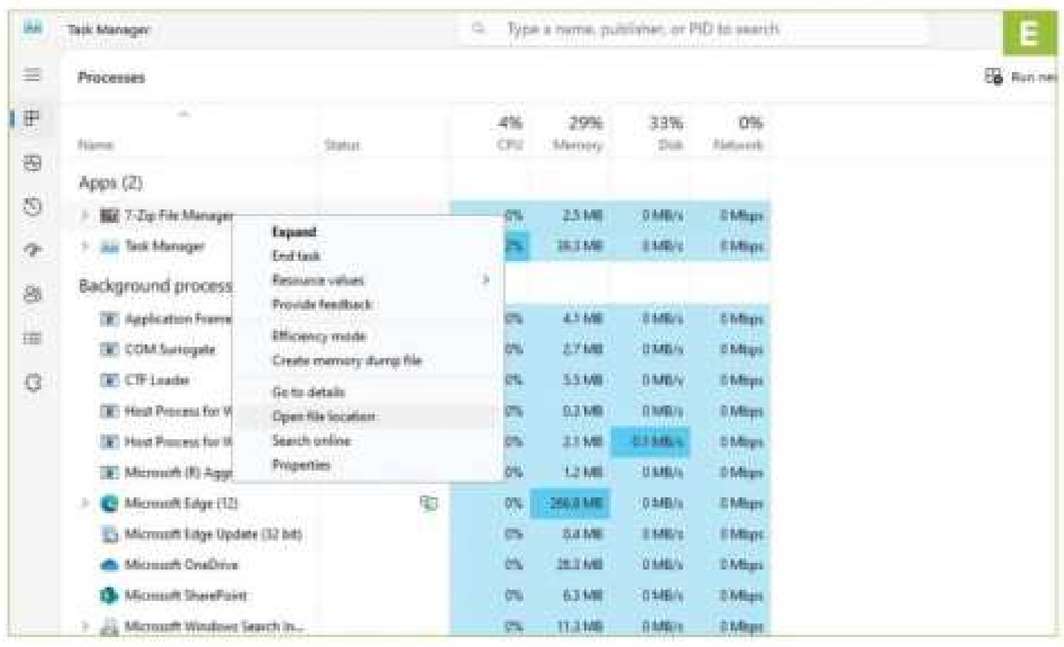

» If the files are harmful, repeat Step 1 to boot to safe mode and rename the folder. We recommend doing this, rather than deleting the files, in case they are actually needed by a legitimate system process.

3 REMOVE AND REINFORCE

At this stage, it's unlikely that any malicious programs remain on your system. If you have a third-party antivirus program installed, feel free to update and run a full scan.

» Even if the infection has now been removed, it could still linger in System Restore points.

» Enter 'restore' into the Windows search bar to select 'Create a restore point'. Once 'System Properties' launches, choose 'Configure', then in the new window, 'Delete' all system restore points for your main drive.

» Browsers are one of the most common methods for stalkerware and other types of malware to enter your system. Now that your browser settings have been reset, you can safely protect it from future infections.

» Launch your chosen browser. You may be asked to sign in to your online account once again. This is usually safe to do, as most browsers don't sync temporary files.

» One of the best ways of preventing malware from gaining a foothold in your system is to install an ad-blocker. Unfortunately, Chromium-based browsers like Edge and Chrome have recently implemented Manifest V3, which limits the effectiveness of these extensions.

» Open-source derivatives of Chromium like Brave, and non-Chromium browsers like Firefox, are unaffected, so this may be a good time to make a switch.

» Whichever browser you use, head to https://ublockorigin.com. This extension will prevent most annoying ads, as well as more harmful code, from loading in the first place. [Image F] We also recommend Disconnect (https://disconnect.me) to prevent the loading of tracking cookies.

» Close your browser and return to 'Windows Security'. If your device was infected, the stalkerware may have disabled key settings. Switch on 'App & Browser Control'. If 'Device Security' reveals that memory integrity is disabled, choose 'Go to settings' to turn this on once again.

» Return to 'Virus & threat protection' > 'Scan options' and run a 'Full scan' to remove any lingering files. While this is running, select 'Firewall & network protection' from the left-hand pane to double check that the firewall is enabled for all network types.

Boost your gaming VPN

YOU'LL NEED THIS

AN ACTIVE SUBSCRIPTION
to your VPN provider

IF YOU'RE A SERIOUS GAMER, you'll be aware of the advantages of using a VPN. As traffic between your device and the server is encrypted, your ISP can't monitor and 'throttle' it, slowing down your game. With the IP address hidden behind the VPN server, players are also better protected from DOS (Denial of Service) and 'swatting' attacks.

Still, connecting to the internet through a VPN introduces an additional 'hop' in your traffic's route, which can increase latency. This can be an issue if your chosen title only allows you split-seconds to shoot, slash, or jump. Many VPN providers claim to offer the very best servers in optimal locations, with swathes of bandwidth for the perfect gaming experience.

In this guide, we'll explore some ways to test such claims. This is done by closely examining the location and type of VPN servers, evaluating different VPN protocols, and deploying extra features and hardware. **–NATE DRAKE**

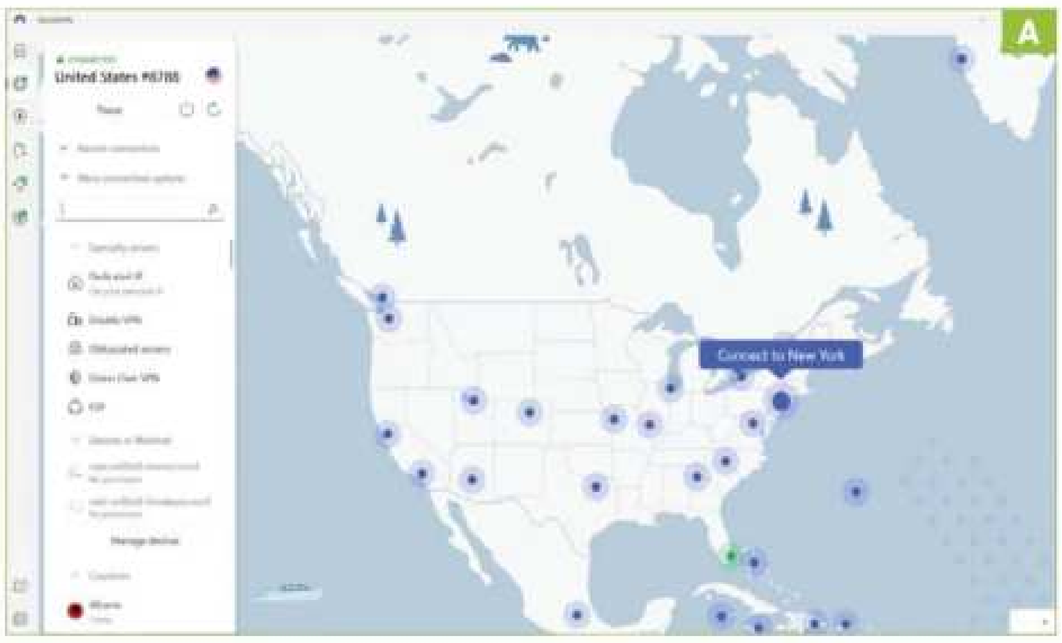

1 SELECT YOUR SERVER

When playing online, every gamer's main consideration is ping time: the time taken between a player's action and the game server's response.

» A good gaming VPN should be able to keep your ping time under 50 milliseconds. If you're engaging in competitive gaming, a ping time under 20 ms is ideal.

» You can keep this down by choosing a VPN server location close to the gaming server. This isn't always as simple as it sounds. For instance, *Fortnite* has dedicated tournament servers in Ashburn, Virginia. If your VPN provider simply offers a 'US Server', this could be located in Moloka'i, Hawaii, which is almost 4,800 miles away. [Image A]

» Even if your provider appears to offer servers close to your gaming location, you should double-check that these are physical rather than virtual.

» Virtual VPN servers have their uses. They can make your device appear to be in a different location, eg. for watching geo-restricted streaming video on Netflix via an appropriate IP address. However, the server on which they run could be located far away from their apparent location.

» This is why you'll experience the lowest latency through connecting to a physical VPN server. It offers shorter distances for traffic routing if it is located closer to your gaming server. Physical servers also tend to be better at allocating shared resources like bandwidth, boosting VPN gaming performance.

» Some titles, like *Fortnite*, allow you to view your ping time via the game network settings. Others, like OpenArena, can display this data via on the leaderboard during play. [Image B] Either way, make sure you measure your ping time both when connected and disconnected from the VPN.

» Alternatively if you know the IP address or URL of your gaming server, open the Windows command prompt and run:

```
ping <url> -n 10
```

eg.

```
ping ping-nae.ds.on.epicgames.com -n 10
```

» This will send 10 data packets to the gaming server, then provide a readout of average, minimum, and maximum latency in milliseconds.

2 PROTOCOL PERFORMANCE

Most VPN servers support connections via a number of different protocols, each of which has their own advantages and drawbacks. [Image C]

» If you're looking for a quick answer on what's best to choose, WireGuard offers very fast performance, as it operates in the server 'kernel space', and only takes up around 4,000 lines of code. It also doesn't compromise on security, given that it uses ChaCha20 encryption.

» OpenVPN is a more established protocol, but has traditionally lagged behind WireGuard in terms of speed. However, if your provider supports OpenVPN DCO (Data Channel Offload), this also operates in the kernel space, offering comparable speeds to WireGuard. If you do opt for OpenVPN, we recommend UDP over TCP, as it usually allows for faster packet transmission.

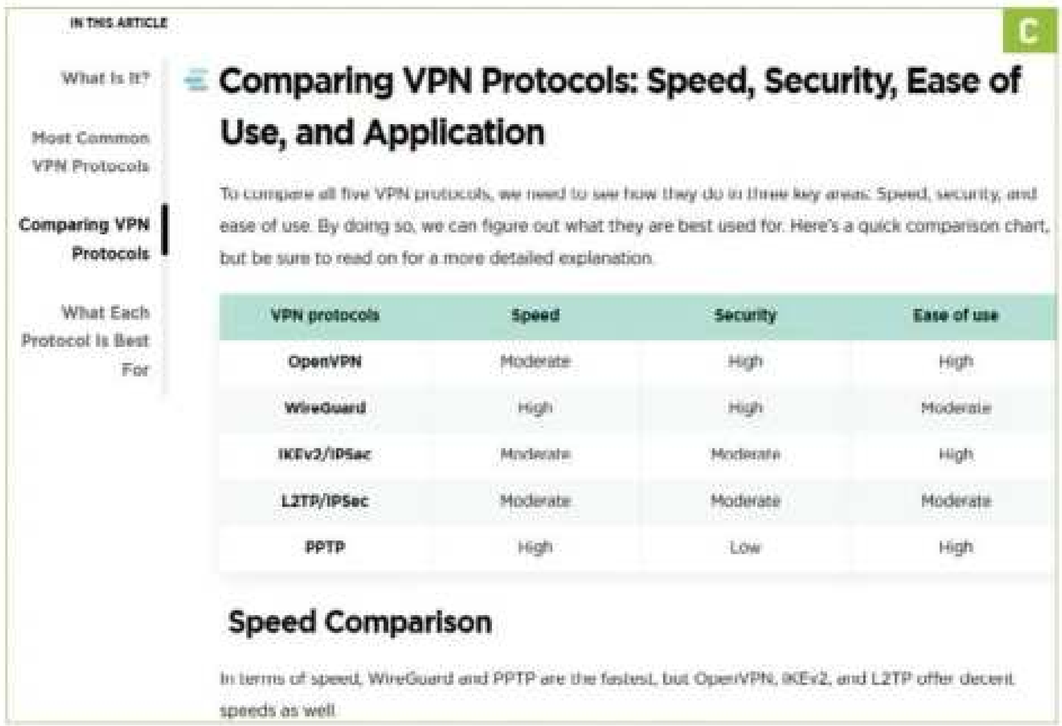

C

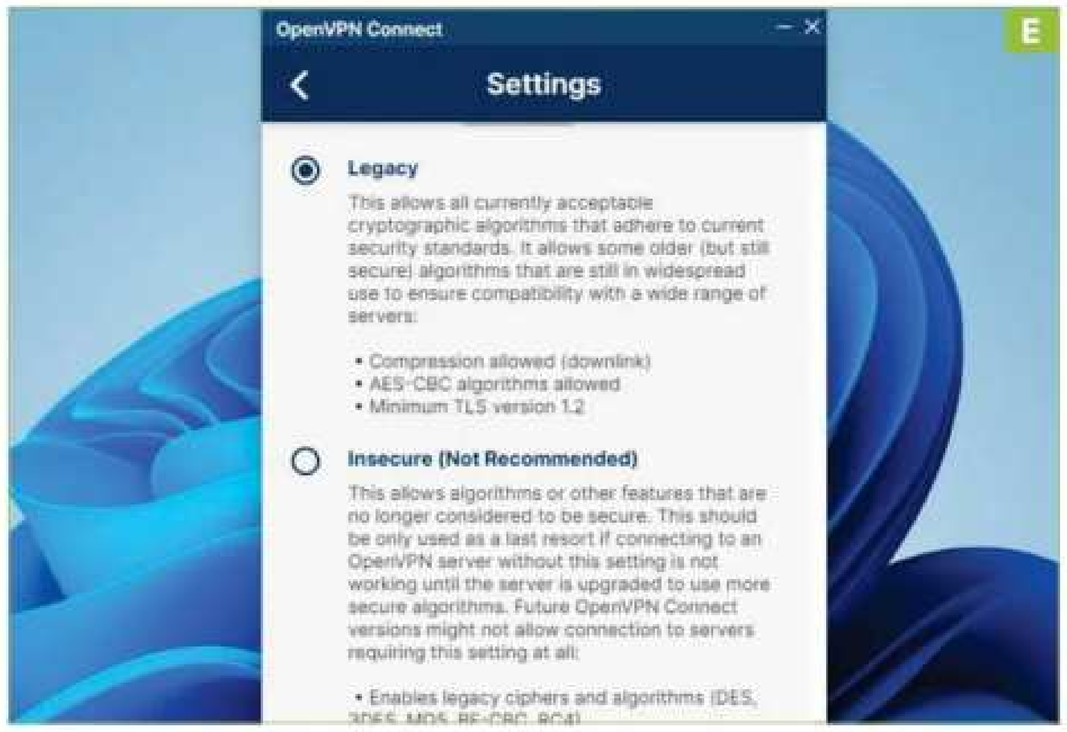

E

» If you plan to game across multiple devices and/or locations, you may prefer to use IKEv2, which has features to handle network changes like dropped connections or switching from Wi-Fi to 5G.

» Gamers are sometimes tempted to use older VPN protocols like PPTP, given that there are fewer encryption overheads. Windows also supports it natively. While this will undoubtedly result in faster gaming performance, Microsoft's implementation of the protocol has serious security vulnerabilities. [Image D] It's also fairly easy to detect and block.

» The simplest way to check and change your VPN protocol is via the client software offered by your provider. Not all servers offer every protocol, so once you've switched, make sure you're still connected to your server of choice.

» While you're using the VPN client, make sure you check through Advanced Settings to see if there are options to improve performance. For instance, TunnelBear offers 'TCP Override' to switch to OpenVPN TCP if UDP is underperforming. This can happen if your ISP or network manager blocks UDP traffic.

» Certain VPN servers may also support using weaker encryption, eg. AES-128-GCM instead of AES-256-GCM, or even legacy ciphers like RC4. [Image E]

» This will generally improve performance, but lower your connection security. Contact your provider to check if this is supported, and use at your own risk.

3 DEPLOY VPN FEATURES

It stands to reason that you'll get the very best VPN gaming performance if you don't use any other internet-enabled apps while playing.

» If this is unavoidable, check if your VPN provider supports 'split tunneling'. This feature is found in certain client software, and allows you to specify which apps will connect to the VPN server, and which will connect directly to the internet. If you select only your game software, this will reduce VPN server load, and can improve performance.

» Don't be afraid to ask your VPN provider about what gaming-specific features they offer. For instance, users of the OpenVPN Connect client can enable the aforementioned DCO mode for faster performance, as long as it is supported by the providers' servers.

» NordVPN's 'Meshnet' tool is ideal for playing LAN Games across the internet. As counterintuitive as this sounds, it establishes secure, direct connections between devices so that they all function as if they were part of a private local network. This can obviate the need for a gaming server for certain titles. Crucially, Meshnet is free to use, even if you don't have a NordVPN subscription. [Image F]

» If you have multiple gaming devices in your home, or simply want the best performance from your PC, you should also consider setting up a VPN router.

» These are routers that are capable of connecting to a VPN server. Naturally, this means that any devices connected to the router's wireless network will automatically use the VPN. This sidesteps the requirement to install and configure VPN client software on your device, freeing up system resources.

» VPN routers are also useful for gaming on devices that don't natively support VPN clients, like the Nintendo Switch. For optimal performance, a VPN router should run community-developed firmware like OpenWrt or DD-WRT. The easiest way to get started is to buy a VPN router with the firmware preinstalled from your VPN provider or online retailer.

» If you prefer to 'flash' the router yourself, first check the OpenWrt website (https://openwrt.org/supported_devices) for a list of compatible devices.

D

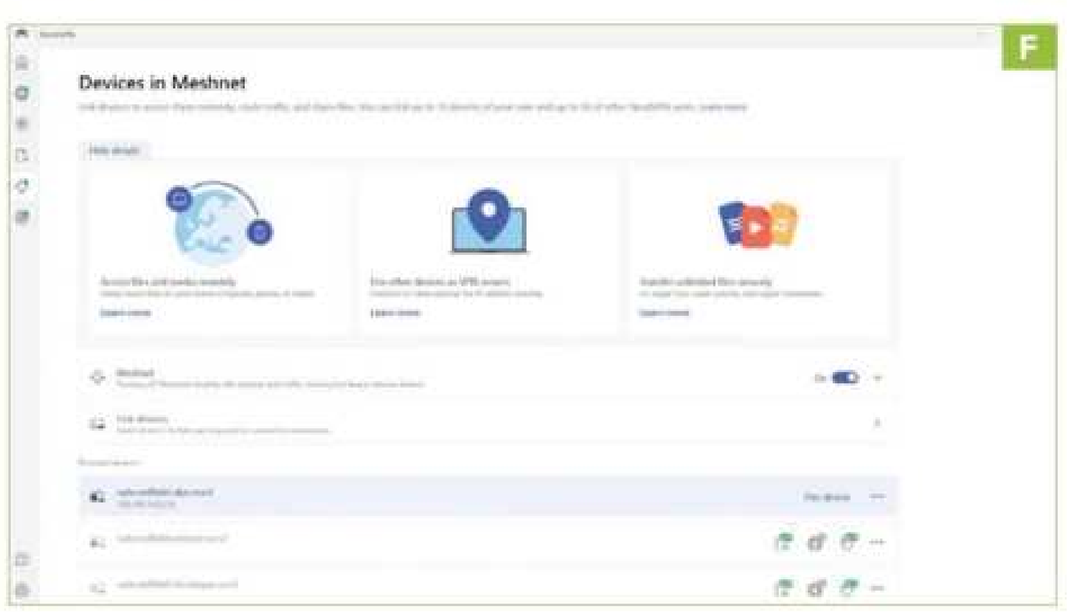

F

How to repair your graphics cards

YOU'LL NEED THIS

HWMONITOR
www.cpuid.com/softwares/hwmonitor.html

3D MARK
https://store.steampowered.com/app/223850/3DMark

WE'VE ALL BEEN THERE. Just picture it for a moment: you're rampaging through your favorite game, decimating monsters left right and center, silver sword drawn, coated in oil, and then suddenly, pop! System crash. What the heck?! The PC reboots, you get back in-game, then five minutes later, the same again. Try one more time. Three minutes later, wham! The game closes to desktop. Frustrated, annoyed, and disgruntled, you shut the machine down. Time to sit down and actually do some work.

We're, of course, talking about the pitfalls and pain of graphics card failure. It's one of the few components that can tell you it's hurting before it keels, and we've become experts at repairing our trusty graphical boats. In this tutorial, we're going to run you through our top tips, covering everything from GPU diagnostics to physical repair work, and software side tweaks that should help alleviate some—if not all—of your Nvidia, AMD, and Intel problems. **–ZAK STOREY**

1 DIAGNOSIS PROGRAMS

First and foremost, we need to identify what is causing the issue. If you've got a relatively new graphics card, it's unlikely to be a physical problem (although it is possible). The first thing we're going to do is look at exactly what's causing the error. Now, you can't typically find these in Window's Event Viewer—it may tell you if a program crashes, but it won't tell you why. Some games do come with error codes on failure, which can be a massive help in diagnosing an issue like this, but that's not a guarantee either.

» So, the first thing we want to do is stress test the graphics card. To do that, we're going to need two programs. First, some form of monitoring software, so ideally you want something like HWMonitor or HWInfo. Secondly, you'll want a benchmark that predominantly stresses the GPU, above all else. FurMark is a great pick for this—however, some errors do only occur with CPU interaction, and as such, we prefer 3D Mark. It's got a number of benchmarks that stress different elements of your GPU so you can identify exactly what's causing the issue.

2 LOOPED TESTS

With your system monitor running, and only the GPU stats up and visible, you're going to want to run these benchmarks on a loop. In the case of 3D Mark, select the benchmark you want to run (Speedway for DX12 and RTX, or Steel Nomad for just DX12), then hit 'Custom Run', and select the looped option. While you're here, you should also disable 'Fullscreen' [Image A].

» If you have a second screen, and can keep HWMonitor on that, perfect. If not, don't worry about it; we'll just run HWMonitor on top. Once the benchmark begins, place your system monitor on top of that. Now for the boring bit: just watch it.

» Ideally, what we want to see is artifacting—things that shouldn't be in that benchmarking. This could be missing textures, purple or odd colors, or screen stuttering and tearing (outside of usual frame-rate drops), or something similar. Also, keep a close eye on your maximum clock speeds and GPU temperature. If this is insanely high, then this indicates that there may be a problem with either the cooling to the card, clearance, or the thermal paste physically on the card itself. If there's major artifacting, but clock speeds and temperatures are within parameters (Google is definitely your friend here, as this will vary, depending on the GPU in question), then we can put it down to being potentially a driver issue instead.

» This is where tech media reviews of your graphics card—if not, then your own database and knowledge—can come in handy. In some cases, for instance, when GPU memory fails, performance will fall off a cliff, but if you don't know where your baseline is to begin with, then user benchmarks can be an incredibly powerful tool to diagnose. Don't worry if you don't have a similar-spec CPU; just check the 4K test results, as most games still don't leverage much CPU power when it comes to that resolution.

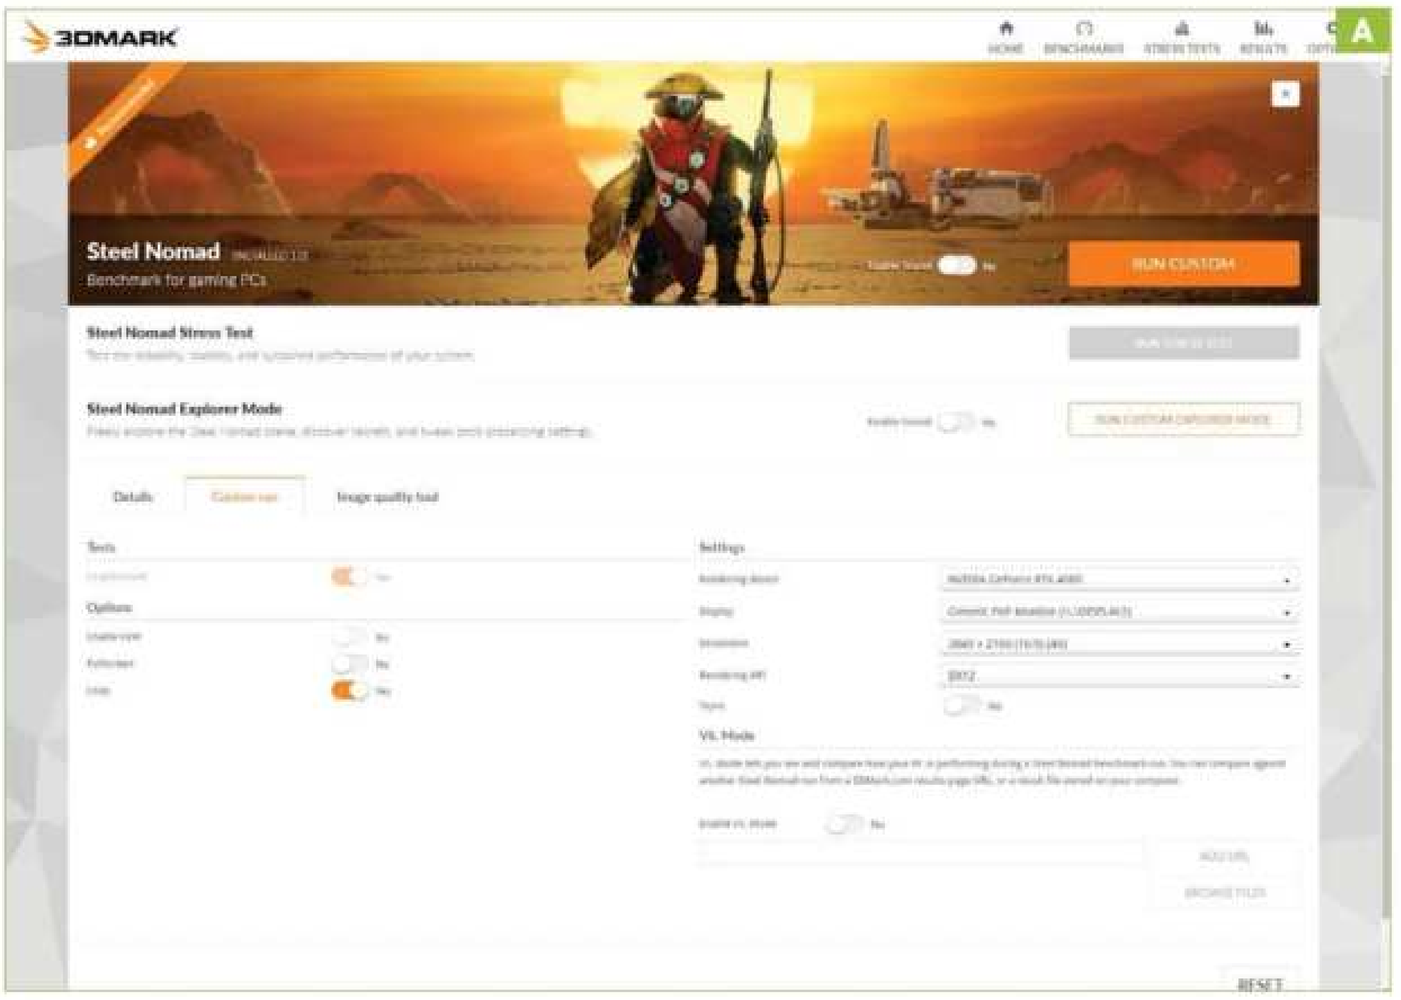

© 3D MARK

3 REPAIRING THE DRIVER

If you believe that it's a driver issue, then there are a number of things we can do to remedy the issue. Don't worry, we're not about to send you to some dodgy third-party driver removal tool. Instead, go to Google, type in either AMD or Nvidia GPU driver, and download the latest version. Both manufacturers' driver installers come with the option to perform a clean install [Image B] during the installation. Select Custom (Advanced) install, then choose which item you want to install, and make sure you check the box that says 'Perform a clean installation'. What this does is remove any and all trace of the old drivers, then install completely fresh drivers on top. It's best to download the driver first, then disconnect your system from the internet to stop Windows automatically attempting to install an older version of your driver as well.

» If you continue to have problems, reinstall Windows in its entirety. Over time, Windows' registry can become filled and with junk and temp files, and possible compatibility conflicts and issues occur. We often recommend reinstalling Windows at least once every six months for optimum performance, if not sooner.

» In both scenarios, once you've reinstalled your driver, you can re-test the system using the same programs.

4 POWERFUL COMPONENTS

The number of times we've had GPU problems, and it's purely arisen from the card not being sat in its PCIe slot correctly, or the PCIe power or 12VHPWR cable not being secure, is beyond count. Even occasionally, PSUs fail on a certain rail, and we need to swap the cables. Always double check that your cables are secure and correctly positioned, as this can cause additional headaches when diagnosing.

» When it comes to speccing your system, or even just upgrading your graphics card, always be mindful of how much power your system may potentially draw from the wall, and how much your PSU can provide. There's a number of wattage calculators online—our favorite is embedded in PC Part Picker. Ensure that you have a PSU with at least 20 percent head room—above what PC Part Picker suggests you will be pulling from the wall—and you'll be just fine. However, be careful, as some of the more budget options, even though they may be rated for 450W, may actually be capable of drawing less.

5 THERMAL DIAGNOSIS

But what if you have an older graphics card? What if it's five or six years old, second hand, or been ran for many moons, night and day, mining in the cryptocurrency boom of the early 2020s? In that case, it may be that the thermal paste connecting your graphics card's GPU to its heatsink has cracked. Yep, this paste can and will dry out from time to time, no longer acting as a sufficient thermal medium between the two components. In that case, you're going to need to carefully disassemble your graphics card.

» To do this, you're going to require a few things. First, a clean work surface, then a small screwdriver set, potentially a 4mm hex nut socket wrench (for some Nvidia GPUs), a fresh tube of thermal paste, and possibly some thermal pads for your memory chips. You can buy the bulk of this from Amazon—make sure you get the right-sized thermal pads for your GPU. You can find that out by speaking directly to the manufacturer, or via a quick Google search.

» When disassembling your graphics card, always start from the rear. The backplate will have a number of screws. Take the ones off around the edges first, then leave the four larger screws in the middle until last (these are the ones that act similarly to a CPU block's screws, securing the GPU heatsink to the GPU). Once those are removed, remove the backplate from the card.

» Take a look at the rear I/O plate. You'll likely find a few screws here as well. Identify which ones connect the plate to the graphics card heatsink, and remove them too, storing them away for later. Finally, remove those four GPU screws. At this point, you should be able to remove the heatsink—carefully pull the two components apart. There should be a few connectors and cables for RGB lighting and fans. These are easy to unclip or unseat from their headers, but be careful when doing this. If the GPU still doesn't unseat, you can go one step further, and gently twist the card itself away from the heatsink. Doing that should loosen the thermal paste, and the two components should fall apart [Image C].

» Now, pay close attention to where the thermal pads are located—you'll need to replace these later if they've frayed apart. Clean off the cracked thermal paste from the GPU itself, and then replace it with the fresh stuff.

» Because this is such a lengthy process, and you're unlikely to see it, we highly advise adding more than what's required. If you've ever installed a GPU waterblock, you'll know that most manufacturers recommend a star pattern of thermal paste on the GPU. Ensure your thermal paste is non conductive, and go to town, before replacing the heatsink and backplate.

» We need to emphasize here that it is very easy to knock off a capacitor, or damage a VRM or Power Phase when doing this, so please do careful, take your time, and do this at your own risk. This will void the warranty on the GPU, but for older graphics cards, it may very well save you some cash.

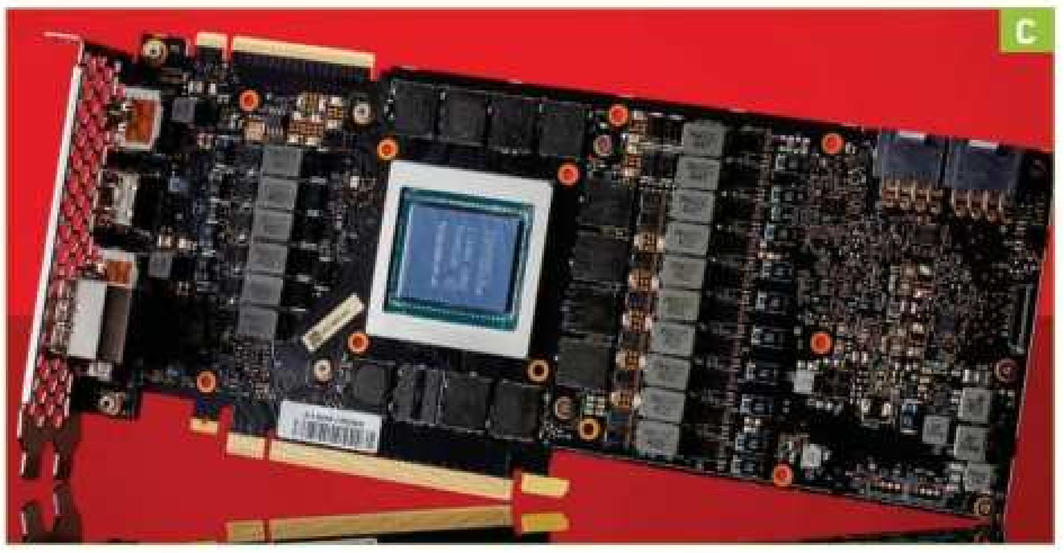
C

LAB NOTES

ZAK STOREY, CONTRIBUTOR

Nvidia's strangle-hold on AI

Taking advantage of UL Procyon's AI Benchmark suite

ARTIFICIAL INTELLIGENCE has gone from buzzword to fascinating topic. Whether that's in the form of data analysis, processes to make workflows more efficient, or debating with LLMs pretending to be Socrates, the ramifications are clear.

I'd heard how much of Nvidia's hardware stock is being pushed into AI development. I needed to understand why Nvidia had an edge. Thanks to the team at UL Solutions, I got access to their AI Computer Vision and Image Generation benchmarks.

It is terrifying how much of an advantage Nvidia has. There's a price difference of around 96 percent (if you pick up a 4080 Super, not our OC variant), and a performance delta of around 93 percent (with ray tracing) in stock testing. But in AI processing, it's a whole other battle.

Take the Windows Machine Learning tests using the GPU. The 4080 OC scores 1,945 versus 1,084 on the 7800 XT—about 79 percent more. Enable those Tensor cores, though, utilizing Nvidia TensorRT, and that difference increases to 148 percent.

The RTX 4080 can, with TensorRT, generate an image in 1.833 seconds. It completed the test in 29.335s, and scored 3,408 points. The RX 7800 XT? Using ONNX, it generated an image in 16.368s, completed the test in 261.885s, and scored just 381. That makes the RTX 4080 nearly nine times faster than its AMD counterpart.

Going back to inference testing, and the contrast is even more stark. Pretty much every chip I've tested fails to break past a score of 56. At best, the NPU in the Asus Zenbook Duo managed 285, and again compared to Nvidia's 2,690 it's just nothing. That worries me, because we've seen time and time again that brands rest on their laurels if they're not challenged.

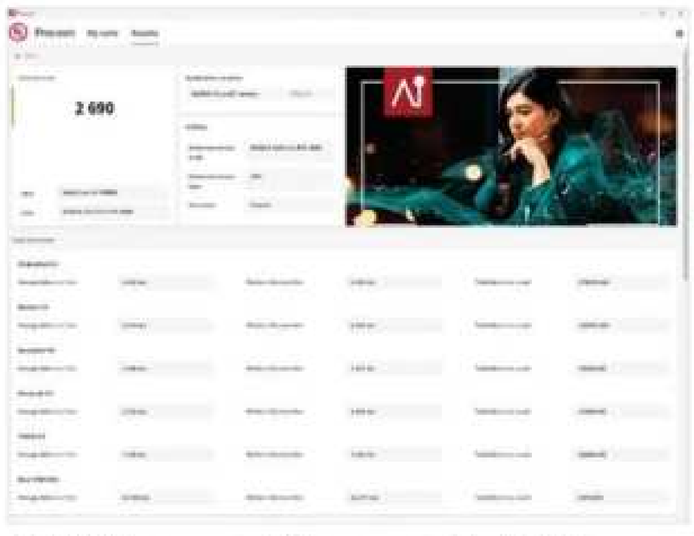

If AI is the name of the game, it's Nvidia.

If we want AI to become the juggernaut it's capable of being, not only do we need to legislate for that across the planet, we also need manufacturers to rethink the direction of their own products. Otherwise, it may very well be that we end up looking at yet another AI winter.

JEREMY LAIRD
Contributor

Just when you thought 32-inch 4K monitors based on Samsung's QD-OLED panel tech were as good as it gets, LG has dropped a bomb. The UltraGear 32GS95UE uses LG's latest WOLED panel technology. Put simply, it's better than the QD-OLED competition.

On paper, it's similar. The 32-inch panel size, 4K native res, 240Hz refresh, and 0.03 ms response performance are dead ringers for the QD-OLED gang. But in a few key areas, the LG is just that bit better.

For starters, it's brighter at 275 nits full screen versus 250 nits. Next, it doesn't exhibit the same warm color balance of those 4K QD-OLED monitors. Plus, the panel doesn't turn gray in bright ambient light.

These matter when you're talking about monitors that tend to cost $1,000 or more. The LG also rocks a Dual Mode feature that uses pixel doubling to offer an alternative 1080p mode, and hits 480Hz. It's better than traditional interpolation, and gives you that sky-high refresh.

As for catches, they are twofold. First, LG has gone for a matte rather than glossy coating. Normally, I prefer glossy on an OLED panel for its heightened sense of contrast. But here, it's just fine. The matte coating is perhaps a little 'glossier' than usual, and somehow it all just works.

The issue is price. It's $1,400—$200 more than most QD-OLED alternatives, and $500 more than the MSI MAG 321UPX. That's hard to justify.

Stats abound. If you like numbers, smart watches are for you.

Editor's Pick:

Samsung Galaxy Watch7

How smart watches can make you healthier

IT MIGHT NOT surprise you that I love stats and data, particularly tied to measurable metrics that I can control and analyze. Fitness is a big passion of mine, and right now, I'm on a bit of a calisthenics hype. This is body weight training, pull-ups, push-ups, dips—things of that ilk. To get stronger in calisthenics, you can do one of two things: reduce your overall body fat levels, or increase muscle mass. Because of that, I like to keep an eye on as many metrics and stats as I can: steps, body composition, stress levels, heartrate. As such, I've advocated for smart watches since their conception.

I had a Samsung Watch Active2 since 2021, decided it was about time to jump ship to a new one, and the Watch7 caught my eye. In fact, I was waiting for it, given I already had the S24 Ultra, and boy is it a beauty, certainly compared to the Watch Active2. A lot of the core specs are admittedly the same or similar; they're both ECG certified, have touch displays, 164 feet of water resistance, and similar AMOLED displays (the Active2 has a higher pixel density, ironically), but the big changes are the OS—the Watch7 featured Android Wear OS 5—and of course, internal storage and CPU performance, with the Watch7 packing in a penta-core 3nm Exynos W1000 versus the dual-core Exynos 9110 10nm chip on the Active2. The battery's larger, too.

I knew going in that regardless of what happened, the Watch7 was going to be better. I did have a choice, though, and the option for the Watch6 was there, too.

At a physical level, they look very similar, bar a more rugged aluminum frame for the Watch7—even the battery is the same. However, it's under the hood where the 7 has the edge, as it features an upgrade from dual-core to pentacore (going from 5nm to 3nm), and has twice as much internal storage (although the same amount of RAM). Reportedly, the GPS connection is far more stable, too, as it's now dual channel, and there are a number of new fitness metrics it can track, along with 'AI' generated health tips and responses to messages, and a 'body battery' metric where the watch looks at your health markers, exercise routines, and sleep patterns, and tries to gauge how much energy you have left in the tank.

One thing that has impressed me is how fast it charges—mind you, it does feature 10W wireless charging. I know there have been reports about poor battery life, but I tend to charge my watch every night before I sleep. Samsung has released an arsenal of patches since its launch.

The fitness routines are mildly baffling for me. I've never liked the idea of trying to track certain exercises like pull-ups or bicep dumbbell curls—it just feels awkward, and that's still the case. However the SPO2, the step counting, body composition analysis, and heart-rate/ECG tech, is all top-tier, and that alone makes it an appetizing pick.

Why the Watch7 over the Ultra? Simply put, there's not enough of a difference between the two to warrant the latter. Despite the Watch7 looking sleeker in my eyes, the internal hardware is near identical beside a slightly larger battery. On top of that, all the sensors are pretty much identical, with the only major advantages being a more ruggedized endurance. That would be fine, but the Ultra's near double the cost of the Watch7. No, just no. **–ZS**

$300, www.samsung.com

© NVIDIA

Reviewed...

The latest Zen 5 architecture is here at last. But does it move the game on for AMD?

AMD Ryzen 7 9700X

The new Zen 5 CPUs are here—time to benchmark!

IT'S BEEN A BUMPY ROAD getting here, what with AMD recalling all stocks of the new Ryzen chips weeks before they were due to be sold, and all about jumping around motherboards, updating firmware, and endlessly fiddling about with settings. But at last, we can bring a review of AMD's latest and greatest desktop CPU design—and it's a little underwhelming.

AMD has stuck to the same fundamental floorplan for the new Ryzen 7 9700X as the last-generation 7700X and even the Zen 3-powered 5700X. Underneath the heatspreader, you've got two chiplets: one CCD (Core Complex Die), which houses all the processing cores and cache, and one IOD (Input/Output Die) that's home to a tiny integrated GPU, PCIe and USB hubs, and the RAM controllers.

The 9700X comes with a fully enabled eight-core chiplet and 5.5GHz top Turbo speed. That 8.3 billion transistor CCD sports the latest AMD Zen 5 architecture, of course, with more L1 cache, more internal bandwidth, superior floating point support, a fancier branch prediction unit, and so on. The changes seem pretty comprehensive and AMD claims that Zen 5 has an average IPC (instruction per clock) uplift of 16 percent over Zen 4, though not every application is going to see such an increase.

The Ryzen 7 9700X also runs a 65W TDP (88W maximum power limit). That's much lower than the 105W TDP of the Zen 4 7700X. It's also half that of Intel's Core i5 14600K and Core i7 14700K. In testing, I couldn't get it to go over 75 C, no matter what task I threw at the processor. With an Asus ROG Strix LC III 360 AIO liquid cooler, it happily ran at 65 C in game testing in a hot office. Impressive.

Anyway, the net result of all those changes is not actually terribly dramatic. Take *Baldur's Gate 3*. Average frame rates for the 9700X are up to 94fps from the 88fps of the 7700X. *Metro Exodus: Enhanced* increases from 147fps to 151fps. Some other titles do show larger boosts in percentage terms. *Homeworld 3* improves from 61fps to 69fps, but these are not gains you're going to actually feel.

It's a similar story for productivity. The Cinebench 2024 multi-thread result only increases from 1,104 points to 1,169, while blender improves from 85 samples per minute to 92. Granted, any performance improvement is welcome, but if these were the kind of gains that we could generally expect from AMD every two years, it would take an awfully long time to see performance increases that you could really feel.

Architecturally, then, Zen 5 is thus far a bit of a disappointment. That could have been offset, had AMD thrown in some more cores. However, it has kept core counts the same across the new model range compared with the outgoing Ryzen 7000 family. What you're left with is largely the same performance as before, just at a much higher real-world price.

Yes, it is true that these chips are launching at lower MSRPs than the last generation. However, Ryzen 7000 pricing has fallen substantially since launch, so you're going to have to pay much more for a 9700X than current 7700X pricing. In pure performance terms, you're getting very little in return. Right now, you can get a Ryzen 7 7800X3D, the 3D V-cache model and best gaming CPU, for only slightly more than this non 3D V-cache representative of the new range. The old 7800X3D flies in terms of frame rates.

This is the problem for AMD when it comes to PC gaming: all such enthusiasts now only care about the X3D models and for good reason. So, if you want a new Zen 5 chip just for gaming, you're better off skipping the 9700X and waiting for the 9800X3D. For everyone else, the lower power consumption and operating temperatures are nice enough, but are they worth the money? **-NICK EVANSON**

VERDICT **7**

AMD Ryzen 7 9700X

+ RYZEN AGAIN Lower power consumption and temperatures; Lower MSRP than before.

– FALLING FLAT Barely any faster than last gen; More expensive real-world prices.

$359, www.amd.com

BENCHMARKS

	AMD Ryzen 7 9700X	AMD Ryzen 7 7700X	Intel Core i7 14700K
Launch price (MSRP)	$359	$399	$409
Cyberpunk 2077 (afps)	110	107	**117**
Baldur's Gate 3 (afps)	94	88	**103**
Homeworld 3 (afps)	**69**	61	62
Metro Exodus Enhanced (afps)	**151**	147	141
Cinebench 2024 multi-thread (points)	1169	1104	**1854**
Blender 4.2.0 (junkshop)	92	85	**132**
CPU measures in gaming (Avg CPU Temp °C)	**62**	69	73

Best scores are in bold.

SPECIFICATIONS

Cores	8
Threads	16
Base clock	3.8GHz
Boost clock	5.5GHz
L3 Cache	32MB
L2 Cache	8MB
Unlocked	Yes
Max PCIe lanes	24
Memory support (up to)	DDR5-5800
Processor Base Power (W)	65
Maximum Package Power (W)	88

The highest-performing RTX 4060 laptop we've ever tested.

Lenovo Legion Pro 5i 16 Gen 9

Excellent performance, but measly memory and storage

THE LENOVO LEGION Pro 7i Gen 8 delivered a superb mix of performance, size, and price. We expected similar results from the Legion Pro 5i Gen 9. Instead, we're feeling miffed by the new Gen 9 model. With impressive gaming performance, you'd think we'd be all over this thing, and yet we can't get over its lackluster memory configuration and low storage capacity.

First, the good stuff—the Pro 5i Gen 9 comes equipped with the Core i7 14650HX. A 24-thread processor, divided into eight P-cores (16 threads with Hyper-Threading) and eight E-cores. That's a lot of processing power for what is ostensibly a gaming laptop on the more affordable end of the market.

The GPU is very much a known quantity: the Nvidia GeForce RTX 4060 mobile, with 8GB of GDDR6 and rated to a total graphics power (TGP) of 140W. That last bit is important, as it determines how much head room the GPU has available, and has a pretty significant impact on the overall gaming performance of the machine. The good news here is that 140W is as good as it gets for the RTX 4060 mobile—a 115W TGP with an extra 25W boost—and it's far and away the fastest RTX 4060 we've tested to date as a result.

Now, let's talk about a few things the Pro 5i Gen 9 doesn't do well. It comes equipped with only a single stick of DDR5 memory, rated at 5,600 MT/s. While retaining 16GB capacity—suitable for a laptop at this price point—a single stick limits the overall bandwidth. It could be much higher with dual-channel memory. While a single stick doesn't appear to hinder the laptop hugely in many games I've tested, it could, it might, and it probably will at some point.

The good news is you can easily access the laptop's innards through 10 Phillips head screws. Adding a second stick or upgrading to a new kit—the latter likely preferable to avoid compatibility woes, but running a higher cost—is as simple as removing the metal shield, popping the old stuff out, and slotting the new stuff in.

You'll likely be in the undercarriage of this laptop at some point shortly after buying it anyway, because the SSD in the Pro 5i is a measly 512GB. When combined with a Windows 11 install and a few Lenovo pre-installed applications (including McAfee), it is a tight squeeze for a few major triple-A games. We ran into issues with storage space during benchmarking, and had to uninstall one game to test another.

Thankfully, there's a spare SSD slot. The primary SSD is one of the faster drives in our testing, so it'd make for a fantastic boot drive alongside a larger data drive in that second NVMe slot. Our issue is we shouldn't have to be ripping the rear off a shiny new gaming laptop to fix what is a pretty poor configuration.

What's frustrating is that this is otherwise a fantastic laptop. The screen is a 16-inch IPS with a 165Hz refresh rate and 2560 x 1600, 16:10 resolution. The keyboard is a full-size unit with a numpad and full RGB backlighting. There are tons of ports, a decent webcam, and a responsive trackpad. Gaming performance is also awesome. The extra oomph from the graphics wattage drives higher performance, and is handled well thanks to its high-quality interior.

There are the usual gaming laptop complaints—battery life is terrible, and it's heavy. There are also competing laptops offering more for less, like the HP Victus with a Core i7 13700H, 16GB RAM, and an RTX 4070, which is $1,199. There's also Lenovo's own Legion Slim 5 with a Ryzen 7, 16GB of admittedly single-channel RAM, and an RTX 4070 for $1,220.

There's a lot to like, but you're going to want to upgrade that RAM and SSD, or spec it up at purchase, which is only going to add to the cost. **–JACOB RIDLEY**

VERDICT 6

Lenovo Legion Pro 5i 16 Gen 9

+ NEXT GENERATION Fastest RTX 4060 mobile laptop; High-wattage GPU; Good screen.

– PICARD Single-channel memory; 512GB SSD; Short battery life; RTX 4070 laptops at similar price.

$1,520 (1TB model tested), www.lenovo.com

SPECIFICATIONS

CPU	Intel Core i7 14650HX
GPU	Nvidia GeForce RTX 4060 mobile (140W)
RAM	16GB DDR5-5600 (1 x 16GB)
SSD	512GB PCIe 4.0 NVMe (2x NVMe slots total)
Screen	16-inch IPS w/G-Sync
Resolution	2560 x 1600
Refresh rate	165Hz
Dimensions	14.3 x 10.3 x 1 inches
Weight	5.5lbs

© UNSPLASH

Asus ProArt PX13

Because professional laptops don't have to be boring

CAN A LAPTOP BE CUTE? This model from Asus certainly makes an attempt, behaving like a manga hero character, as it's all chibi on the outside, but reveals great strength when put to the test.

The ProArt PX13 is a 13-inch two-in-one machine, so needs to be comfortably graspable when in tablet mode. This means it can't taper to the sort of edge that might make the machine useful as a weapon in the event of a zombie invasion, and results in a slab-sided look with widely rounded corners. This is a good thing, as not only does it allow you to pick it up and fold the keyboard back, so you can pretend you're in *Star Trek* while all you're actually doing is compiling a grocery list, but with a stylus in the box, you can even do some drawing, too.

The screen is a lovely 3K OLED touchscreen that's only let down—and it's not a big problem—by a maximum refresh rate of 60Hz. Why is this not a problem, given that we expect our games to run at 4K/240Hz these days? The answer is the full-size HDMI 2.1 port on the edge of the laptop, which means you can hook it up to all kinds of external displays and really let the RTX 4070 loose. A laptop like this will spend much of its life in such a position, as it's aimed squarely at those who prefer to create rather than consume.

Backing this assertion up is the presence of both the stylus and what Asus calls the 'Asus Dial' on the trackpad. It's a circle, essentially, and hooks into some customization software to provide a slightly different way of interacting with Adobe's Creative Cloud apps and others. Twirl your finger over the right spot, and you can scrub back and forth on a timeline, or expand or contract a brush tip. It's not as good as a physical dial, which we've seen on Asus ProArt laptops in the past, but it does take up less space, and is useful once you get the hang of it. You can also just use it as a volume control in Windows, if you want.

Being an OLED two-in-one, putting the machine into tent mode and using it to watch movies is a definite possibility. It's as good in this configuration as it is in its other shapes, and the screen is bright and colorful. You might even forget that 60Hz refresh rate. Hook up a Bluetooth controller, and it makes for a decent gaming machine, too.

The most interesting thing about the ProArt is probably the Ryzen AI 9 processor, which brings an NPU plus four full AMD Zen 5 cores and eight of the cache-limited Zen 5c units with it. Those CPU cores are all capable of multithreading, too, which gives them more potential power than Intel's single-thread efficient cores. The default TDP is 28W, below Intel's Ultra 9 185H, and can boost up to 5.1GHz. In tests, it's up there with the i9s of the 13th and 14th generations, and beats the Core Ultra 7. Using the RTX 4070 means endurance will take a hit—this shows in the battery life test, where the PX13 managed three hours and 34 minutes of constant use.

Asus may have gone out to make a laptop that's aimed at creative users, but it has built a great machine for the rest of us, too. If you want to game after a hard day creating art in Photoshop or Blender, it's got the chips for you. **–IAN EVENDEN**

VERDICT 9

Asus ProArt PX13

PROART OLED, Zen5, RTX 4070, and a load of other good stuff.

LOWART You'll be wanting an external monitor to go higher than 60Hz.

From $1,999, www.asus.com

SPECIFICATIONS

CPU	AMD Ryzen AI 9 HX 370
Graphics	Nvidia GeForce RTX 4070
Memory	32GB DDR5
Screen size	13.3-inch OLED
Resolution	2,880 x 1,800
Refresh rate	60Hz
Color coverage (stated)	100% P3
Storage	2TB SSD, microSD card
Connectivity	Wi-Fi 7, Bluetooth 5.4, 2x USB 4.0, 1x USB 3.2 Gen 2 Type-A, 1x HDMI 2.1, 3.5mm audio
Dimensions	11.74 x 8.26 x 0.7 inches
Weight	3lbs

© UNSPLASH

The Asus Dial sits at the top of the trackpad, and is easily accessible.

Gigabyte Aorus CO49DQ

One of the best ultrawides, if you can stretch to it

WE'VE ALL DONE IT when buying a new PC: Let the imagination run wild and picture not only the case, studded with RGB fans and radiators, but the entire setup. The precise mouse we want, the keyboard with its mechanical switches, perhaps a headset positioned artfully nearby. Then there's the monitor. No one goes into PC building hoping to push together two widescreen displays that they only bought because they were cheap. What we all want, deep down, is one of these.

It rises on an elegant single foot from the desk, the cables connecting it to your PC (or PCs, since this is a fantasy) neatly hidden. There is no dividing line where two screens come together; instead your Windows desktop is smoothly continuous across the space two monitors would take up. Switch it on, and it is both dazzlingly bright and beautifully saturated, games leap from its surface, and the curve brings the edges neatly into your peripheral vision. For the first time, you're happy.

The market is now full of OLED ultrawides, although this fact doesn't seem to be bringing their prices down. This model from Gigabyte has a good balance of features, with enough inputs (including USB-C that can send video and data signals while charging your PC, too) to satisfy almost everyone. Picture-in-picture and side-by-side modes offer versatility, especially if you give up the HDMI port to a streaming stick or game console (in our testing, the HDMI couldn't accept the screen's full resolution, so is better relegated to alternate purposes), but they're packed a little too close together, which can make inserting and removing cables tricky. And with just five inches of vertical adjustment on offer,

you may want to consider an arm, though check it can hold the weight first.

We sit closer to computer monitors than we do to TVs, which means the benefits of both the twice-1440p resolution and the way OLED displays colors and contrast are more visible. Add to this the curve—a flat ultrawide monitor requires you to turn your head to see what's going on, but the curve brings things into your peripheral vision, so you can detect movement. You may still need to turn to address flanking soldiers, but the payoff is worth it if you can get your chair in the right position.

The downside of this is size. The Aorus CO49DQ arrives in a box the dimensions of a small couch that will both annoy the FedEx operative who delivers it and your neighbor when you ask for help getting it up the stairs. It comes in an enormous amount of packaging, but that's probably essential to keep the screen safe.

So, if you're thinking of buying a matched pair of 1440p monitors to go with your PC, consider this instead. The lack of a dividing line down the middle is a game-changer, and OLEDs are the best way to experience PC gaming. The ability to sit in the sweet spot and have the action wrap around you does wonders for immersion. It's not cheap, but then good things rarely are. **–IAN EVENDEN**

VERDICT **8**

Gigabyte Aorus CO49DQ

BROAD Wider than wide with a useful USB hub.

FLAWED A little expensive, and has plenty of competition.

$1,100, www.gigabyte.com

SPECIFICATIONS

Panel	49-inch QD-OLED, 5,120x1,440, 144Hz
I/O	1x DisplayPort 1.4, 2x HDMI 2.1, 1x USB-C (18W PD), USB hub/KVM, 3.5mm audio
Color response (claimed)	99% DCI-P3
Brightness (claimed)	250 nits SDR, 400 nits HDR
Speakers	2x 5W
Tilt	-5° to +20°
Height adjustment	5 inches
Dimensions	1,195 x 559 x 261.5mm, 47 x 22 x 10.3 inches
Weight	25lb

© UNSPLASH

The support foot is sturdy, but the amount of height adjustment is low.

Helios's signature RGB light bar makes a return here.

Acer Predator Helios 18

Perhaps you don't need a desktop PC anymore

THERE'S BEEN A RACE to make laptops smaller for years now, with extreme thinness and lightness prized in portable machines. Acer has had enough of that, and its Predator Helios 18 is about as big as a mobile PC can get in 2024.

Actually, is it mobile? This isn't a laptop you're going to want to throw in a bag and take with you, or have on your lap for very long. It's heavy and gets hot, and while its components are absolutely ideal for any kind of rendering work, whether that's in Alan Wake 2 or After Effects, its natural habitat is firmly on top of a traditional desk, rather than casually laid on a bed or balanced in the back seat of a car.

This is the triumphant return of the desktop replacement, a category of laptop we thought had been lost, with mere 16-inchers the best we could look forward to. The great benefit of a machine like this is how tidy it can make your working or gaming space—the important cables (power, USB-C, and HDMI) all slot into the back rather than spread out of the sides, and while you can't charge it over USB-C (it takes a massive 330W external brick to power it), you can use a single USB connector for video and peripherals if you have a monitor with a hub. Then you can fold it away and lock it in your desk, adding a layer of security to the tidiness.

Once it's up and running, the combination of 14th-gen i9 and RTX 4090 is top-end. It laughs at ultra settings and makes chicken noises at ray tracing. You're looking at frame rates that can outstrip the 250Hz mini-LED IPS screen, and getting hold of a large, fast, external 4K OLED monitor should be up on your list if you have the space and budget.

There are, of course, drawbacks to a machine such as this. It's very expensive—up there in MacBook Pro territory. It's hot, which means it's loud. We caught it spinning its fans a bit when in sleep mode with the screen closed, and when called upon to perform ray tracing, made a rushing sound that summoned someone from the floor below to see what the noise was. That said, the cooling system, which uses two fans and liquid metal thermal grease, is clearly effective, as despite being pushed to 100% in a ray-tracing demo, the GPU didn't rise above 68 C, even though the CPU was up at 87 C.

Battery life is also an issue. While it's possible to get more than four hours out of the Helios 18, no one's buying this to scroll websites or make video calls. Fire up that big GPU and you'll be lucky to get more than two hours of use out of it.

Alongside gaming, we can see this being popular among the video-editing and 3D-rendering crowd. Turn off the RGB lights and the dreadful sound it makes at startup, and you've got something that will munch through a barrage of rendering and AI tasks. Photoshop's generative fill and Lightroom's denoising tasks are completed in a flash.

Give it a firm enough place to stand, and the Helios 18 will move the Earth. Just because it's a laptop doesn't mean you can carry it around. This chunky specimen shows that desktop replacements never really went away. **–IAN EVENDEN**

VERDICT **9**

Acer Predator Helios 18

+ HELIOS Excellent performance and a large screen.

– ALBATROSS Heavy, hot, and hungry for battery power.

$3,100, www.acer.com

SPECIFICATIONS

CPU	Intel Core i9 14900HX
Graphics	Nvidia GeForce RTX 4090 (16GB)
Memory	32GB DDR5
Screen size	18-inch IPS
Resolution	2,560x1,600
Refresh rate	250Hz
Storage	2TB SSD, microSD
Connectivity	Wi-Fi 7, Bluetooth, Gigabit Ethernet, HDMI, 1x USB 3.2 Gen 1 Type-A, 2x USB 3.2 Gen 2 Type-A, 2x USB 3.2 Gen 2 Type-C
Dimensions	1.14 x 16 x 12.2 inches
Weight	7.17lb

© UNSPLASH

Hyte Thicc Q60

Almost more mobile phone than CPU cooler

TAKE A LOOK at the Hyte Thicc Q60, and you might come to the conclusion that this is an all-in-one liquid cooler that pushes appearance over performance. That's not strictly true, however. Yes, the Q60 is adorned with a 5-inch screen, a litany of RGB LEDs, and some tidy cabling, but there's also a chunky CPU cooler beneath it all.

The Q60 is a big beast. It comes with a radiator double as thick as some, coming in at 52mm. That's reportedly to use space that's not often utilized to its fullest in modern PC cases, according to Hyte—it's opting for a thicker radiator in the place of a longer one.

The Q60 is billed as a 240mm radiator on account of the two 120mm Thicc FP12 fans included, but it's actually 288mm from top to bottom. The remaining space is taken up by the pump, which takes residence on the radiator rather than in the more traditional spot behind the cold plate. The benefit of this, for the Q60 at least, is more room for the enormous screen above the CPU.

The 5-inch screen on the Q60 is best described as a tiny computer. It includes a quad-core Arm Cortex CPU, with 2GB of DDR4 running at 2,666 MT/s, and running at 1.3GHz. It also has 32GB of eMMC storage, all to run the many widgets, which Hyte has named 'faces', available via the Nexus application, which is pretty much a must-install with this cooler.

I've opted for a parade of system-monitoring 'faces': liquid in temp, liquid out temp, CPU temp, and GPU temp. Oh, and the clock, which I cannot seem to remove. There are a lot of options to choose from, and a few more 'coming soon'—clock speeds, fan speeds, and throughput are all reportedly on their way. You can also set the screen to display your media playing via Spotify, and play your own media. Plus, you're free to change the background.

We're not usually fans of sticking a screen on peripherals and parts for no discernable reason except 'we could'. But we're sort of into the Q60's central screen. It's impressive, and not at all as domineering as we had originally expected. The subtle two-tone design of the rest of the cooler helps—if the fans were also adorned with RGB lighting, it would probably feel a bit too much. As it stands, the screen is pretty wicked.

Yet, as we mentioned at the beginning, the Q60 isn't only about appearances; it is a capable liquid cooler in its own right. In balanced mode, with the PC idling, the Q60 is nominally audible. During gaming, it's only a little noisier—we've been running the *Metro Exodus* benchmark many times over, and the only noise coming from the Q60 throughout is a steady, polite hum. It doesn't noticeably fluctuate between speeds, either, which makes the noise it does produce much less distracting.

Only during more CPU-intensive workloads did we begin to take note of the sound coming from the fans and pump unit. In return, this cooler delivered the lowest temperatures while gaming out of those we tested, which included a 360mm AIO. The Q60 isn't as capable as others when it's pushed to the limit, however.

There are other reasons to consider the Q60. It's extremely neat and tidy, for one. It requires only a handful of cables, collected into a single dual-USB Type-C connector, and plugged seamlessly into the radiator. The fans can also be bought separately and daisy-chained together with a dedicated cable-free connector or specially fitted cables. It also comes with a six-year warranty, which matches the likes of Corsair and Arctic.

The Hyte Thicc Q60 is smart, good-looking, and quiet. It delivers exceptional cooling performance in games, too, though trips up in more demanding thermal tests. It's by no means essential, but it embodies everything great about custom PC building by working well and looking cool while doing it. **–JACOB RIDLEY**

VERDICT **7**

Hyte Thicc Q60

HIGHEST HYTES Quiet; Stunning; Impressive cooling performance for gaming; Handy for system monitoring; Zero-RPM modex.

LOWEST LOWS Less impressive with demanding CPU workloads; High idle temperatures; Expensive.

$300, www.hyte.com

SPECIFICATIONS

Compatibility	LGA 1700, 1200, 115X, 2011*, 2066* - AMD Socket AM5, AM4, TR4* (*available from Hyte support)
Dimensions	120 x 288 x 52mm (cold plate: 56 x 56 x 1.5 mm)
Radiator	288mm, aluminum
Pump	Ceramic, up to 4,500 RPM
Fans	2x Thicc FP12 120 mm, Fluid Dynamic Bearing, up to 3,000 RPM
Lighting	Full RGB on radiator logo, array behind screen
Screen	5-inch IPS, 60 Hz, 300 nits, 720 x 1280 resolution
Storage	32GB eMMC
Memory	2GB DDR4

The 5-inch display shows all the system info you need.

Bag a software bargain

Don't pay full price for software when we can offer you huge reductions on everything from security software to Windows 11 Pro

Norton 360 Premium for $29.99

WE'VE NEGOTIATED a killer deal with Norton. No subscriptions; just a one-off bargain price of $29.99, compared to the regular $109 charge. That buys you 15 months of cover from the Norton Security suite across 10 devices

SAVE 72%

Because it's the Premium version, you get a host of extra tools. There's 75GB of cloud storage for secure backups, plus the full version of Norton Secure VPN, both to protect your identity and provide a way of watching TV from home while abroad.

Norton Parental Control provides tools to see exactly what your kids are up to on their various devices, and you also get GPS location monitoring for Android and iOS.

To round things off, Norton Password Manager generates and stores passwords across all your devices, while SafeCam for PC stops cybercriminals taking photos on your webcam without your knowledge.

ALSO CONSIDER **Avast Ultimate Suite 2024 for $29.99 (10 devices for two years) SAVE 87%**

Windows 11 Pro for $79.99

WANT TO UPGRADE from Windows 11 Home to Windows 11 Pro? Or perhaps you're building a new PC from scratch, or looking to move up from Windows 10? Maybe you want to run Windows on your Mac? Whatever your motives, a Windows 11 Professional OEM license is a great choice.

The first reason is simple: price. Microsoft charges $119.99 for a retail license of Windows 11 Home, and a staggering $219.99 for Windows 11 Professional. By choosing an OEM license—and avoiding the many dubious sources of licenses you may find online—you cut that price substantially.

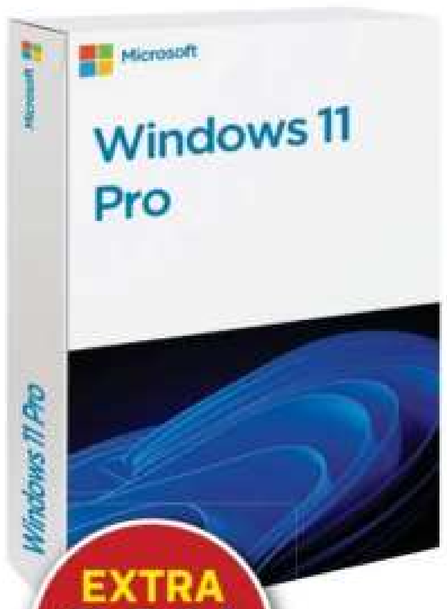

EXTRA 20% OFF!

USE PROOCT20 AT CHECKOUT

ALSO CONSIDER **Windows 11 Home for $49.99 SAVE 64%**

CCleaner 6 Professional for $14.95

CCLEANER HAS A reputation for bringing an old computer back to its best, and we've negotiated a 50 percent discount on the Pro version. This will remove unwanted information, temporary files, browsing history, log files, and even settings from old software.

One new CCleaner 6 feature is its Performance Optimizer. Unlike optimization techniques in other apps, the boost it gives can be measured in benchmarks. It uses a three-step process, detecting active apps, services, background tasks, and processes. Anything it deems unnecessary is put into sleep mode. The result: better performance and battery life on laptops.

Use code PROOCTP20 at checkout for 20 percent off.

EXTRA 20% OFF!

USE PROOCT20 AT CHECKOUT

ALSO CONSIDER **Adobe Photoshop Elements 2023 for $70 SAVE 30%**

software.maximumpc.com

NZXT C1500 Platinum

Top-tier performance and efficiency

THE C1500 PLATINUM, the latest addition to NZXT's power supply lineup, is priced at $370. That alone limits it to a niche audience, along with the few people who actually need that sort of power in 2024. 2025, with its rumored even higher-powered GPUs, could be another matter.

However, if you do find yourself in that camp, you'll be pleased with its exceptional performance and reliability. With features like fully modular black nylon-sleeved cables, a Sunon MagLev 140mm fan, two PCI Express 5.1 12V-2x6, and world-class electrical performance, it's also designed to meet ATX 3.1 standards.

Considering the tier of this PSU, its bundle is surprisingly minimalistic. NZXT provides only the necessary mounting screws, an IEC C19 AC power cable, and a purple bag for the modular cables. The highlight here is the presence of two 600 Watt 12V-2x6 connectors and six PCI Express 6+2 pin cables with just one connector on each, forgoing 'piggybagging' multiple PCI Express connectors on a single cable.

It's housed in a chassis that measures 180mm in length—significantly longer than the standard ATX dimensions specified by the ATX design guide. This length is unsurprising, bearing in mind its high power output, so it requires a case with ample clearance to fit properly. Like other NZXT PSUs, it has a 'Zero-Fan' switch, but the speed of the fan will still be controlled by the unit's thermal control circuitry, so it will be running at minimum speed until the load is high enough to justify an airflow increase.

The NZXT C1500 Platinum PSU is equipped with a Sunon MagLev MFE0251VX-1Q010-S99 140mm fan, featuring a magnetic levitation engine. The patented Sunon engine design offers superior performance, and the fan's performance is not affected by orientation. While MagLev fans can be noisy at higher speeds, they excel in high-temperature environments, with the fan in this model reaching the terrifying maximum speed of 2,500rpm.

In our tests, the NZXT C1500 Platinum PSU easily meets the 80Plus Platinum certification standards, missing the 80Plus Titanium certification by a hair when tested with an input voltage of 115 VAC. When tested with a 115 VAC input, it achieves an average nominal load efficiency of 92.6 percent, which increases to 93.7 percent with a 230 VAC input. The unit's efficiency peaks at around 40 percent of its capacity. Also, the PSU demonstrates impressive efficiency even at low loads, nearing 87 percent with a load of just 70 Watts.

HOT TOPIC

During hot testing, the NZXT C1500 Platinum PSU displays a minor decrease in efficiency under heavy loads, with recorded figures of 92.2 percent at 115 VAC and 93.3 percent at 230 VAC, compared to 92.6 percent and 93.7 percent during cold testing. This reduction is inconsequential, and occurs almost linearly across the load range. There are no signs of thermal stress, even at maximum load. The PSU is rated to perform efficiently at ambient temperatures up to 50 C, and operates seamlessly under these conditions.

All of this means this is an exceptional PSU for performance and efficiency. The 140mm Sunon MagLev fan, equipped with a magnetic levitation engine, ensures superior cooling performance and durability. The NZXT C1500 Platinum PSU excels in protection and build quality. Manufactured by Channel-Well Technology (CWT), a reputable OEM, the unit features top-tier Japanese capacitors from Rubycon, Nichicon, and Nippon Chemi-Con. Its solid construction and high-quality components are backed by a 10-year warranty, providing users with peace of mind. However, its price tag is a drawback. This substantial cost positions it in a niche market, appealing primarily to the few users who require its massive 1500W power output. For many, the performance may not justify the premium price, especially when comparable alternatives exist at lower price points. –EMMANOUIL FYLLADITAKIS

VERDICT **9**

NZXT C1500 Platinum

PLATINUM Very high efficiency; Superior cooling; 10-year warranty.

BRONZE High price; Niche market appeal; Large size.

$370, www.nzxt.com

SPECIFICATIONS

Outputs	1 x 24-pin ATX Power Cable, 1 x 4+4-pin CPU Power Cable, 1 x 8-pin CPU Power Cable, 6 x 6+2-pin PCIe Power Cable, 2 x 16-pin (12+4) 12V-2x6 PCIe Cable, 12 x SATA Cable, 4 x Peripherals Cable
Fan	140 x 140 x 25mm
Noise	41.0 dBA max
Dimensions	5.91 x 3.39 x 7.09 inches
Weight	5.34lbs
Warranty	10 years

BenQ X300G 4K Short Throw Projector

Priced high, yet punchy

THIS IS A BEAUTIFUL piece of kit. The X300G from BenQ sits somewhere in the middle of its latest projector product stack. It's complete with 4K 3LED projection lens, exceptional keystone accuracy, auto-adjustment, and impressive brightness and response times, all packed into a nice, attractive little box. It's advertised as portable, and to a certain extent it kind of is, but it also very much kind of isn't.

We actually looked up the word 'portable' to find a modern definition. Google sums it up as a smaller version of something that can be easily carried. By that logic, the X300G is absolutely that: portable. Don't confuse that with being able to use this anywhere and everywhere, however, at least without an electric hookup. There's a notion in tech that 'portable' means battery powered. Certainly, the XGIMI Halo+ ticks that box, if only for a few hours, but sadly it's not something that the X300G can muster in response. It's dinky, and that's about it.

That is one of the few marks against the X300G. The reality is that this thing is an absolute beast when it comes to projection. It's got an awesome lens, capable of driving 2,000 lumens of brightness at 4K, 1440p or 1080p, with a 30-bit color setup for 1.07 billion colors, and you can tell, because the images from this thing are awesome to look at.

We do a few tests when it comes to projectors: one up against a gray wall (a bit of a torture test), and the other against a white projection screen. On the gray wall, surprisingly this still pops. The darker tone does affect the brightness, but it only makes it look like a standard SDR monitor, and that alone is incredible. Throw up some HDR 4K content, and you're blown away by the color accuracy and clarity, even with a 60-70 inch throw. On a white wall, though, you're in for a hell of a time.

The X300G is in a lot of ways similar to the X3100i: the form factor, a small box, the lens (although that does utilize a 4LED projection system), and even the response times bring it in line. The one area it does struggle in, however, is on the audio front. BenQ has a patented underside 'air radiator' that acts as a miniature subwoofer. Otherwise, the audio is relatively tinny and underwhelming, certainly compared to something like the X500i or X3100i, which both feature punchier audio.

Cooling is pretty good—there's plenty of ventilation on either side, and you can barely hear the fans, which is ideal when you're scrolling through menus. If you are looking to game on it, however, 4K might be something to avoid. As standard, it lands with a 16ms response time, making it sluggish at best, particularly with that 60Hz refresh rate. It's okay for console gaming or less competitive play, but if you're keen to jump into some PvP gaming on a PC, then running at 1440p with 8ms response or 1080p at 4ms is the better way to go. Similarly, it supports 1440p at 120Hz and 1080p at 240Hz, too.

Otherwise, the overall chassis design is impressive. It's small, compact, the I/O is somewhat limited, but bearable, and the Google TV implementation is great. It's even got some stylish little orange LED lighting on it, accentuating the orange highlights glimmering around the lens. The only problem? The price. At $1,700, it's not cheap, and given the X500i is a hundred dollars less, for a slightly better projection lens, sleeker form factor and better audio, it makes the X300G a hard pill to swallow. It's good, just priced incorrectly. **–ZAK STOREY**

VERDICT **8**

BenQ X300G

THE OG Outstanding projection tech; Form factor is adorable; Google TV implementation is solid.

LIKE A G6 Audio is sub-par; No battery; Pricing is a bit skewed.

$1,700, www.benq.com

SPECIFICATIONS

Resolutions	4K@60 Hz, 2K@120Hz, 1080p@240Hz. 60Hz limited to 16.7ms, 120Hz 8.3ms, 240Hz 4.2ms.
Display Colors	1.07 billion
Contrast Ratio	600,000:1
Brightness	2,000 lumens
Light Source Life	20,000-30,000 hours
Zoom Ratio	1.2x
Throw Ratio	0.69 - 0.83
Projection Offset	102.5%
Connectivity	USB 2.0 Type A, HDMI 2.1, USB-C DP
Weight	6.6lbs
Dimensions	8.3 x 7.1 x 7.7 inches

Asus ROG Azoth Extreme

A gaming great with an absurdly heavy price tag

THE ROG AZOTH EXTREME is a $500 gaming keyboard. You're probably reading this, mouth agape, wondering what it makes it worth anywhere near that price. In many ways, that's how I've approached this review, because I didn't get it at first, and to be honest, I still don't.

First impressions are very good. I lifted the lid of the box it arrived in, and it unwrapped itself like a cake in a Wes Anderson movie. Inside the second branded box—truly no expense spared here—sits the Azoth Extreme. I reach in and lift it out, which takes a little bit more effort than I was expecting, as it weighs 3.3 pounds.

The first thing to notice is it's a 75 percent size—the same as its predecessor, the ROG Azoth. The case is hewn out of a block of aluminum, and stained to a deep black. The underside has been preened by a CNC machine down to a futuristic military look, with the outer edges milled away into a stepped pattern. For a slab of milled anodized aluminum, it's one of the loveliest I've laid my eyes on. I set about attaching each of the feet—they're magnetic and snap easily in place.

Beneath the board sits a conspicuous cutout, inside which sits a Type-A USB dongle and a switch marked 'soft' on one side and 'hard' on the other. If I slide the switch to hard mode, the switch feel is firmer, more responsive, and bouncier; if it's in soft mode, it's flatter and less snappy. As I type this on the Azoth Extreme though, I don't actually remember which mode I've set it in.

The package includes a product that Asus hopes to sell separately someday: the Polling Rate Booster. This is an extender for the included Type-A USB dongle that greatly increases the polling rate of the device up to 8,000Hz. That means it communicates with your computer 8,000 times a second, or eight times more than most.

I've tested a high polling rate gaming keyboard before: the Corsair K70 Max. It's not a feature I care much about—I'm not an eSports pro, and I don't notice the difference between 1,000Hz and 8,000Hz. This does highlight one concern: who is its target audience? The high polling rate may appeal to competitive gamers, but they might prefer a keyboard with more features, including analog switches and rapid trigger, like the Wooting 60HE.

IN TESTING

The Azoth Extreme's switches are fantastic for gaming—the linear NX Snow switches inside my review sample are pre-lubed and bliss to game with. Each offers a fast response as I test the board in hard mode, every key press swiftly bouncing back into place as I desperately mash at them in a game of *The Finals*.

What the Azoth Extreme is best for, however, is typing. I spend my days tapping away, and for the past few weeks it's been me and the Azoth Extreme. I come away from that experience feeling that it is a wonderful way to go about my work. It doesn't matter whether it's in hard or soft mode; the light clack and springy response of each key is great for rapid and satisfying typing experience.

The problem, though, is that it's no better than the best tenkeyless gaming keyboard today, the Keychron Q3 Max, which costs $200—less than half the price of this Asus. There's a certain joyous thud to the Q3 Max that wins me over. There are a few other keyboards besides the Q3 Max that are worth considering for pure typing feel, such as the Ducky ProjectD Outlaw65 (reviewed in our March issue), and Asus ROG Strix Scope II 96 Wireless (November 2023).

The one feature you really have to love to justify the purchase is the OLED touchscreen. It does things like capture the waveform of the music I'm listening to on Tidal, and I can flick the screen to switch up the animation. It's neat—often pointless, but neat.

What you're left with is a gloriously high-end gaming keyboard held back by a runaway budget. –JACOB RIDLEY

VERDICT **7**

Asus ROG Azoth Extreme

OLED Typing experience, Useful media controls, Equaliser function, Beautiful chassis.

OH NO High price, Can be matched by cheaper keyboards.

$500, www.asus.com

SPECIFICATIONS

Size	75%
Base	Hot-swappable
Switches	ROG NX Snow/Storm
Backlighting	Per key
Display	Full-color touchscreen OLED
Anti-ghosting	N-key rollover
Polling rate	1,000/8,000Hz (with Polling Rate Booster)
Dimensions	13.1 x 5.5 x 1.6 inches
Weight	3.33lbs (w/out cable/wrist rest)
Connections	2.4GHz, USB Type-C, Bluetooth
Battery life	130 hours (w/out lighting)

World of Goo 2

PHYSICS-BASED PUZZLE

Goo-d enough for two

IT HAS BEEN 16 YEARS since *World of Goo* squidged its way into our hearts and hard drives. The much-loved physics-based puzzle game was one of modern indie gaming's earliest and biggest successes, but a lot has changed since then.

Polished games from unfamiliar studios with barely a dozen staff members aren't some headline-grabbing curiosity anymore. They're absolutely everywhere, all the time. They're often discounted, bundled up, and at times even given away for free. Who needs more *World of Goo*, then, when the original was lovely and creative, and very much done and dusted years ago? Me, actually. I do. We all do, really.

Playing this game is a lot like catching up with an old friend. It looks, feels, and plays in much the same way it did before. Everything from the artistic stage select design to more than a few goo types, environmental hazards, and puzzle pieces call back to the original gooey head-scratcher.

Clearing the first few challenges wasn't just easy; it felt instinctive. There I was, watching towers of goo dangerously sway as I grabbed, built, and stretched an increasingly wobbly mass toward the exit pipe that would suck it all up and end the level. Just like the ancient past of 2008, sometimes I needed to hook balloons onto my goo lattice to move it around, flipping it over and over to 'walk' it across the landscape or encourage it to roll in a particular direction. At others, I needed to set something on fire, or watch out for drafts.

There are a lot of obviously recycled ideas in here, and that's good. These old favorites don't exist in an innovation-free vacuum, but are mixed in with high-speed train rides, little boating expeditions, tense journeys through the dark, and even a spot of goo-themed golf. So whenever an old concept makes a comeback, it feels like the welcome return of a good idea, a greatest hits compilation of some of the finest puzzling PC gaming has ever seen.

Goo 2 isn't rehashing old favorites because it's out of ideas; it wants to show you a good time, while offering a few amusing signposts and cutscenes filled with pointed comments on corporate greenwashing, commercial-minded environmentalism, and the cynicism of sustainable capitalism along the way.

This good time is elevated by generous flexibility. The goo is as pleasantly tactile as ever. It's always satisfying just to imagine something and then try to build it, even if the dangling mound of goo balls you end up with aren't always the most efficient use of your time and resources. Moreover, tough challenges for each stage are always available, and yet entirely optional. You can ignore them, or come back to some of them later on.

Maybe you don't fancy this puzzle right now. Maybe this particular area's puzzle theme doesn't do anything for you. Maybe you just want to open everything up right away and then dive into whatever you feel like taking on today, happily saving the rest for later like the treats they are. It's all up to you.

As for flaws, lots of levels use a hazy wash of something non-specific as their backdrop color, which looks great in screenshots, but means the helpful white

World of Goo 2 is endlessly flexible: you can take on tough challenges or completely ignore them.

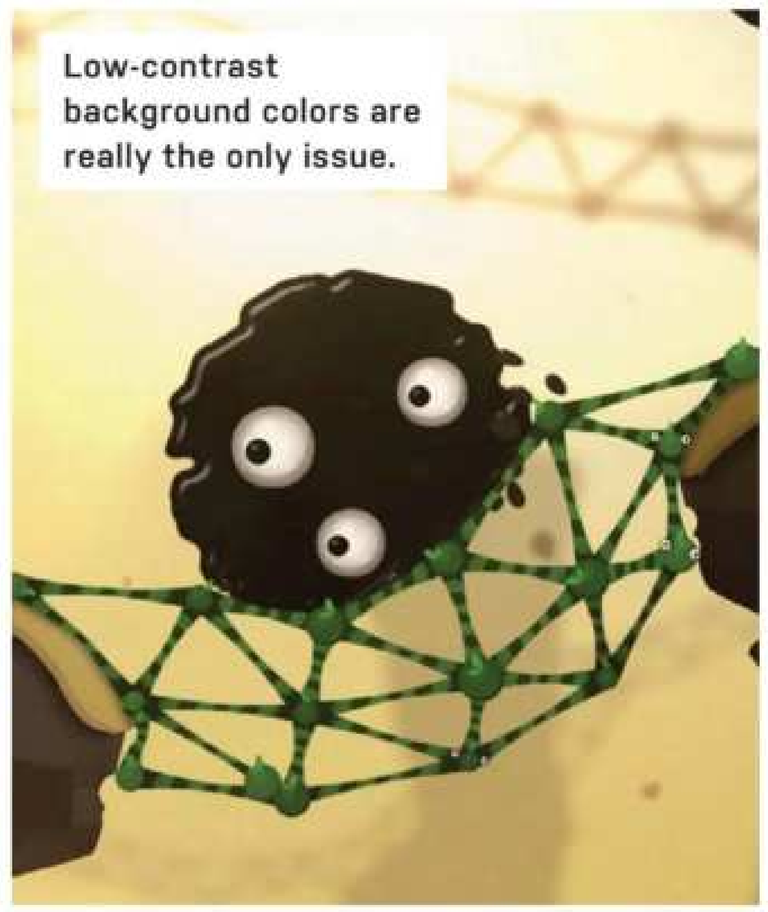

Low-contrast background colors are really the only issue.

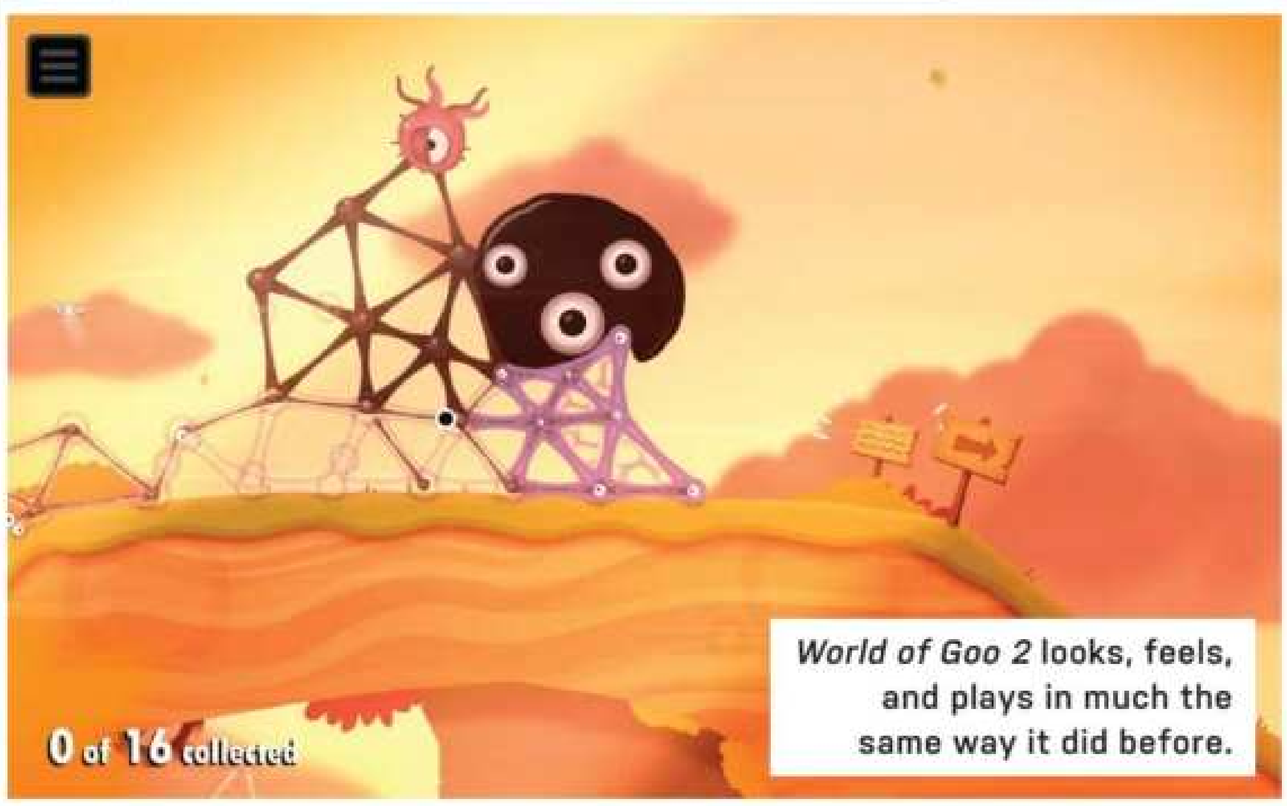

World of Goo 2 looks, feels, and plays in much the same way it did before.

Innovations include high-speed train rides, boating expeditions, and even a spot of goo-themed golf.

lines used to indicate exactly how the floating balls of goo connect are hard to make out. In most levels, this just means it can take longer than necessary to build a solid bridge of goo. In others, it can be a challenge run-ending issue.

But this was honestly the worst thing I found in *Goo 2*, and even when it did cause me problems, I was generally having too much fun to really mind the odd stumble. It's a clever, surprising game that celebrates all the goo that came before and all the goo here and now. It's goo(d) to be back. –KERRY BRUNSKILL

VERDICT **8**

World of Goo 2

GOO(D) TIMES Everything that made the original great remains; Endlessly flexible.

STICKY SITUATION Background colors can be problematical, even challenge-ending.

RECOMMENDED SPECS CPU, Intel Core i5-2500K. GPU, Nvidia GeForce GTX 960. RAM, 8GB. OS, Windows 10.

$29.99, www.worldofgoo2.com, M-rated

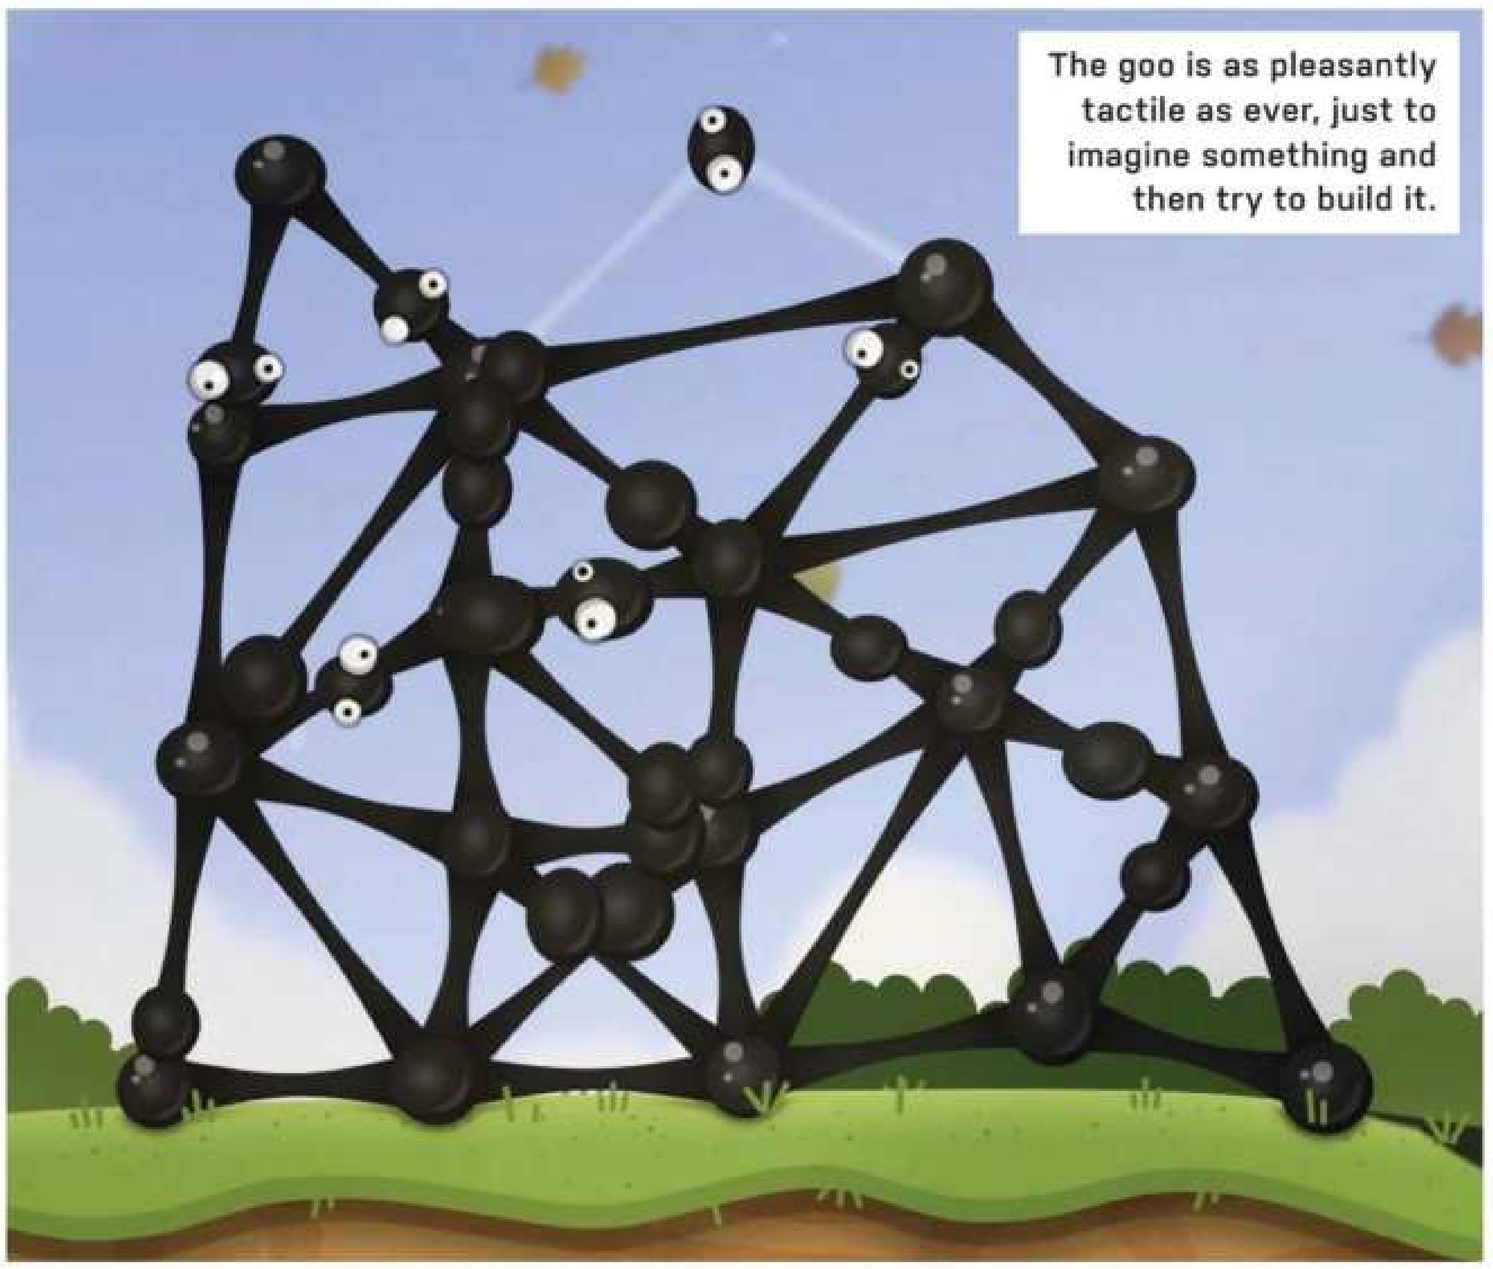

The goo is as pleasantly tactile as ever, just to imagine something and then try to build it.

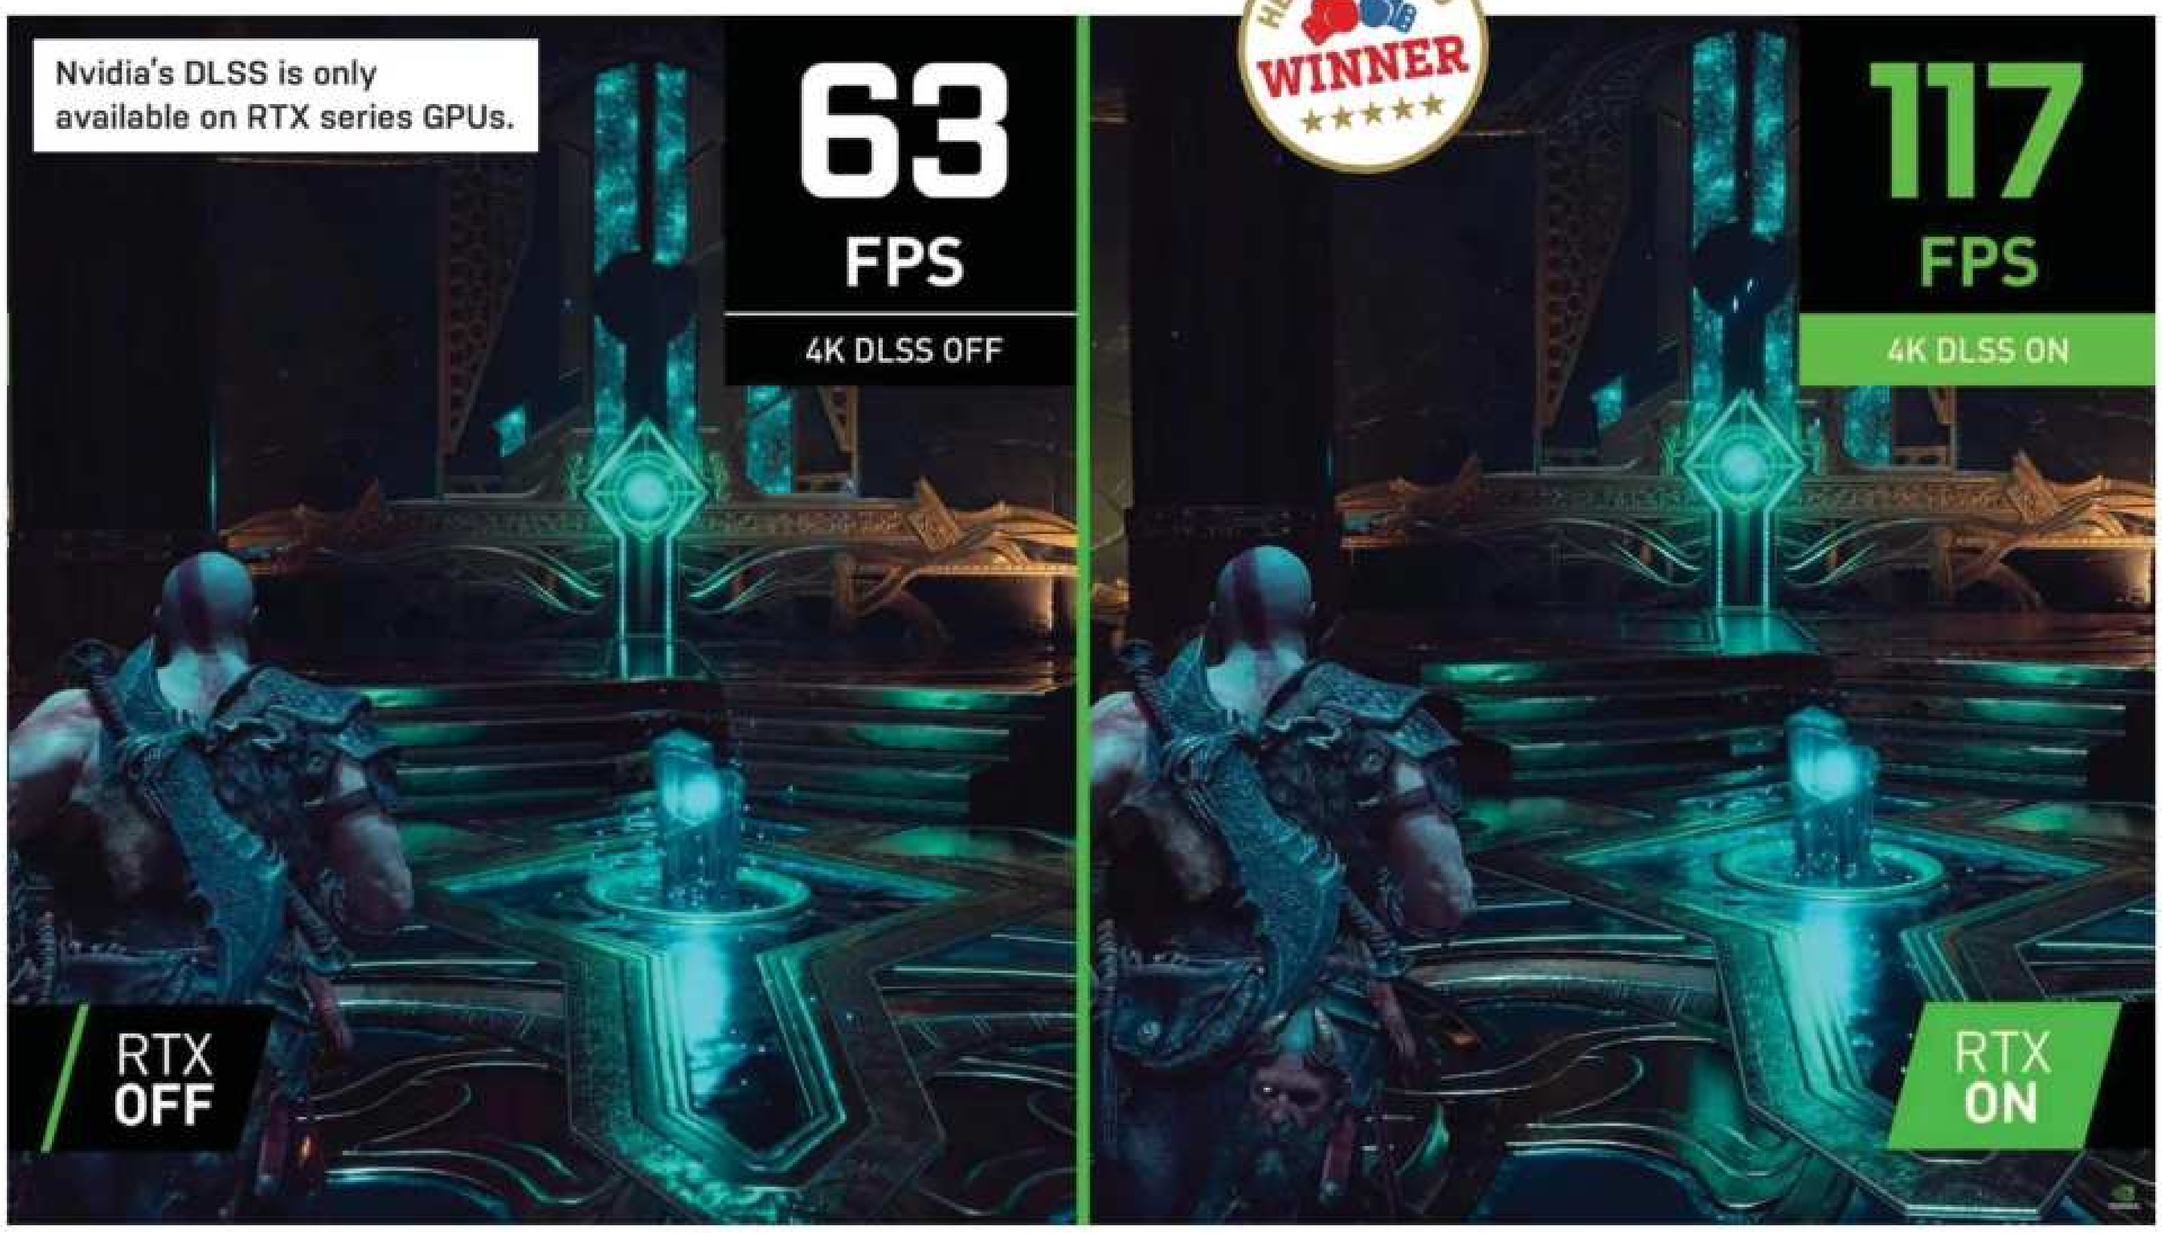

Nvidia DLSS vs AMD FSR

Which AI upscaling technique has the edge?

4K GAMING is a problem. It's always been a challenge to render at that high a resolution. If anything, the res arrived long before any graphics card or console could touch it. The problem is simple: compared to 1080p, 4K represents a 300 percent increase in the number of pixels on display at any given time. We effectively quadrupled the total workload overnight by shifting to a 4K resolution. Not only that, but even in 2024, eliminating jaggies is as important as it ever was. Anti-aliasing efforts similarly are incredibly graphically intensive. Jaggies occur because games render textures in pixels, and because pixels are square, you end up with artifacts creating a jagged texture, often along the edge of a model.

We ended up in a challenging situation where the resolution and refresh rates increased, and although pixel density had shot up, jaggies were still a thing, and the methods we were using to eliminate them via anti-aliasing likewise cost even more performance when processed at 4K.

To get around this, Nvidia began to implement DLSS in its first-generation RTX graphics cards. With the added benefit of Tensor cores in its cards, it could finally leverage machine learning and AI supersampling (predominantly done on its own supercomputers to create models for the tensor cores to utilize) to effectively eliminate that jaggie problem and upscale textures. Initially, DLSS 1.0 focused mainly on anti-aliasing. It was only with DLSS 2.0 that we received the supersampling as well. In that case, instead of rendering a game at 4K by enabling DLSS, the GPU would drop native resolution and upscale that back to 4K, using AI to identify the patterns based on the models provided by Nvidia's own supercomputers, and simultaneously improve frame rates in the process.

DLSS has come a long way since then, and we're now in its 3.5 iteration, complete with Ray Reconstruction and AI Frame Generation, as well improving things further (as long as you have Nvidia's low latency mode enabled, frame-gen does add lag). It's now featured in over 400 titles and growing.

It's not alone in the battle for Super Sampling dominance. AMD launched its own form, known as Fidelity FX Super Resolution, and it operates in a similar manner to how Nvidia runs its own DLSS. However, AMD's variant does not rely on AI Machine Learning models for general upscaling, instead using them to polish some Fluid Motion Frames. That's probably a good thing, as AMD's own AI accelerators found in its latest generation of cards still aren't on par with Nvidia's dedicated Tensor units. But the question lies in which has a competitive edge over the other? Which should you be using?

THE CASE FOR DLSS

On the surface, DLSS feels like a much more rigid AI technique. It's limited to Nvidia cards, it's bound to the hardware itself, and if you're running anything other than an RTX series GPU, you're out of luck. That said, the one major advantage it does seem to have is that when DLSS is updated, it goes live on every game at the same time. In AMD's case, the latest versions of FSR require developers to go back and update their titles with support for it. That is problematic, as not every development team has the capacity or even desire to do that. DLSS now has a massive game database as a result.

On the flipside, FSR is open to any graphics card—you don't necessarily

4K "Extreme" graphics preset with "Extreme" raytracing setting and FSR 2.2 "Performance" mode. See endnote RS-513

need an AMD GPU to run it. In fact, you can run it directly on an Nvidia card, or an Intel GPU. In fact, we've done some testing with exactly that premise in mind.

Utilizing an Nvidia RTX 4080 and Intel Core i9-14900K, we've tested *Cyberpunk 2077* at 4K. Bear in mind that we're using the Ray Tracing Ultra preset, and AMD's FSR 2.1, and DLSS 3.5. Although FSR 3 is not yet implemented in-game, it does support Frame Generation in a similar manner to Nvidia, with its own low latency adjustments embedded in the driver. However, it's worth noting that you can add Frame Generation into FSR 2.1 while using an Nvidia GPU.

In DLSS's Ultra Performance mode, it nets 115.47fps on average (146.98fps with frame gen), on Balanced that drops to 62.88fps (89.98fps with gen), Quality nets 50.93fps (76.69) and with no DLSS you get 26.99fps. The difference is staggering.

That doesn't come without some loss, however. Most of the changes come in the form of eliminating jaggies and improving image clarity. Anything below balanced is somewhat underwhelming. Still frame generation is the biggest net benefit, and doesn't impact quality of the image itself.

FSR UNDER THE MICROSCOPE

In the case of AMD, FSR performs admirably in the scenario, even despite being an older variant. With Ultra Performance, it lands in at 110.97 (145.61 with frame gen), 60.54 on balanced (86.80 with fg), and on quality you get 49.96 (and 74.85 with frame gen).

Generally speaking, the difference isn't massive—DLSS does have an edge on its quality presets, particularly when it comes to AA. There's a number of scenes in the *Cyberpunk* benchmark with grates on the floor and edges of worktops, and in those cases DLSS has an edge, but it's still hard to spot the difference. Nvidia similarly does seem to have an advantage with ray-traced reflections under DLSS over FSR, with anything above 'Quality' looking more blurry under AMD's tech, while the same result only occurred with Nvidia above the 'Balanced' preset.

As you can see, performance between the two is tight. Nvidia does have an edge with DLSS in *Cyberpunk*, but again, we're talking about a new version of its DLSS tech, running in one of its latest GPUs, while FSR is an older version on a non-AMD GPU. The reality is they're likely both close to even in overall performance.

THE WINNER?

Both are phenomenally good at what they do. They represent a categorical change in how we can render graphics in-game, and for the most part, both are free.

Nvidia is a touch more exclusive and locked off—you do need an Nvidia GPU alongside an RTX card, and 3.5 in particular is further locked behind its 4000 series GPUs, while in AMD's case, a broader range of GPUs are compatible with FSR 3 and its Fluid Motion Frames tech. Still, it's not as widely supported, and requires developer intervention, which is where Nvidia has the edge. Although it has Frame Generation too, the number of titles with support is incredibly limited.

If you have an RTX card, we recommend checking out how your game looks on a case-by-case basis, and how each form impacts frame rates. This does vary from title to title. Some games look better and run smoother with FSR 2 over 3, and some are better with DLSS entirely. Still, there's no better time to be a GPU enthusiast, that's for sure. –ZAK STOREY

VERDICT 9

Nvidia DLSS

DEEP THINKING More widely adopted; Easier to deal with updates; Performs incredibly well; AA edge.

DAY DREAMING Hardware locked; Latest tech often requires latest GPU gen; Frame gen leads to latency.

Free, www.nvidia.com

VERDICT 9

AMD FSR

FREE THINKING Compatible with far more cards; Performs just as well; Latency is better.

FATALLY SINKING Requires developer time to implement; Slightly worse AA performance; Reflections not quite as tight.

Free, www.amd.com

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> Irresponsible Intel
> What the BTF?
> AI, AI, Oh?

Unsettled by Intel

I read with intrigue and a small amount of concern in your September issue that you'd decided to drop Intel's 13th and 14th generation CPUs from your Blueprint builds that issue, given the ongoing concerns around their stability.

I was concerned, because I'd literally just used your mid-range blueprint to put together an Intel-based system using your recommended Intel Core i5-14600K processor. The chip is now in my system, and has been running fine, and I'm now outside of the return period where I can take it back to the retailer.

Do you have any advice on what I should be doing? Should I be requesting a new chip from Intel? Should I be asking for a replacement of a 12th-gen chip, which shouldn't be affected? I'm obviously trying to limit the amount of downtime that would result from having to return a CPU, now I finally have a system in place. I'd ideally just go for an AMD equivalent, but that obviously requires effectively rebuilding my PC with a brand new motherboard.

I'm a bit lost, to be honest! Please help!

–R. Thatcher

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS:

As we've noted a lot through this issue and recent editions, while the situation around Intel chips is becoming clearer now thanks to motherboard firmware updates and an extended warranty policy, it still doesn't feel like we're out of the woods yet. There seems to be an inherent hardware problem with these chips, and while there are things that Intel, hardware partners, and users can do about that, there's still a chance that these chips will fail at some point in the future.

Intel may well have an extended warranty in place to mitigate that outcome, but as you've noted, being without a CPU in your main system, even for a matter of days, can be a major inconvenience. It's also not clear to me whether Intel is able to address the issue on replacement hardware going forward—I suspect it won't, so you may well be replacing one faulty chip with another, and then what happens if it fails outside of the warranty period? True, that may be so far off that it's not worth worrying about, but if you do keep your CPUs for around five years or more, then it's concerning.

As someone who also has a i5-14600K in my system, I can tell you what I'm doing—I'm keeping an eye on my motherboard manufacturer's website at least once a week at the moment, and applying whatever new motherboard BIOS they've released, which feels like much more regularly than normal specifically to address the Intel CPU issues.

My motherboard is an NZXT N7 Z790, but if you followed the Blueprint in the August issue, you may have the MSI Z790-S. For now, I'm making sure not to apply any overclocks to the CPU in the BIOS or in Windows, and I'm using the NZXT CAM software to shift the cooling into 'Performance' mode when I'm gaming or doing anything CPU-intensive, just in an effort to keep the processor as cool as possible. I'm probably being overcautious, but like you, I just don't want to have to go through the inconvenience of having to RMA a broken chip.

In the long term, I'll probably move to Intel's 15th or 16th generation chip in the next couple of years, although given Intel's recent issues, I'm now much more likely to consider an AMD equivalent. I do think Intel will learn its lesson and be much more likely to produce stable chips going forward, but reliability is so important for those of us who use our PCs for work, I understand if others just don't think it's worth taking a chance.

BTF? WTF!

I loved reading about your BTF build in the last issue (page 16), but I have to admit, I had to grab my reading glasses to double-check I was reading the price correctly in the ingredients list. $4,821!? I could get a liquid-cooling specialist to build me a machine to the same configuration, and have an almost-silent PC with the

↘ submit your questions to: editor@maximumpc.com

same power for that price! Or I could build one myself, with non-BTF parts, for about half the price! Who do manufacturers seriously expect to be selling this stuff to!?

I get it, it's early days for the BTF form factor, and being able to finally hide away all those cables is a cool concept, especially for those of us who care so much about the aesthetics of our builds. But I'd personally much rather accept the necessity of wiring in my PC, and spend a bit more on nice-looking custom cables, or just go for a PSU that comes with black or white power leads that match the color of my build.

As you noted, buying BTF components at a premium price and then probably seeing those items depreciate faster than standard non-BTF components, just because fewer people can use them when you sell them on, is the nail in the coffin for the BTF concept.

Overall, great idea, but stupid execution, and I just can't see BTF taking off. Great to be able to read about it in the pages of *Maximum PC*, as I'd rather it be you than me spending close to $5k to check this new technology out!

–E. Watts

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS:
Yes, the pricing for the build in the last issue wasn't an error, I promise! It was the BTF premium you had to pay for a compatible motherboard, case, and graphics card, which all added up! Although the use of a high-end CPU, power supply, and two SSDs certainly didn't help things either! However, after spending over $3,000 on just the BTF case, motherboard, and GPU, I imagine our builder Zak was thinking, "In for a penny, in for a pound"

Reader E. Watts queried the price of last issue's BTF build. Yep, no mistake!

when it came to using other premium components!

I think broadly speaking, Zak and I agree with your assessment of BTF at this particular juncture, and you can watch us discuss this in our Story of the Build video linked to on page 17 of the September issue. It is really cool, it looks really tidy, and once you've built a BTF PC, it's hard to go back to the mess of wires usually found in a standard system. But as you say, there's a huge tax to pay as an early adopter, and you're likely to get less for your old BTF components when you come to sell them. There are other issues you didn't mention, too, like the fact that Zak found it quite tricky to actually build a BTF system, despite his decades of experience, not to mention that there are *still* cables to be seen, such as those coming off your CPU cooler. So it's a long way off perfect.

As with all things technology, though, prices will come down, more people will adopt it, and who knows, maybe over time, BTF just becomes the standard at some point down the line? Of course, there's also the chance that it's all a huge misfire, it's dead on arrival, and it becomes a footnote in the history of PC building that we all laugh about in 10 years' time. For now, I'm prepared to give it the benefit of the doubt, although I'm a long way from investing in a BTF system myself.

Rise of the machines

I was wondering when *Maximum PC* would succumb to the allure of AI and feature an Artificial Intelligence-based cover, and I got my answer in the September issue. To be fair to you, Zak's feature was actually really interesting—definitely a *Maximum PC* take on the whole nascent industry, which is what I enjoy about the magazine.

I think my apprehension about the term 'AI' is that it has becoming overused by tech companies eager to sell you an upgrade that you don't really need. It's also often compared to the internet in terms of its potential to revolutionize everything from work to communication, and even the way that we're entertained. I'm old enough to remember the move to the connected world when I first got access to the internet at college, and let me tell you, it was far more exciting than being able to get ChatGPT to summarize a report, or DALL.E mock up an image.

I am, of course, being a little glib, and I realize I'm using AI regularly, just in ways that I've already become accustomed to. Nvidia's DLSS is extraordinary, and has become my default option in games to improve frame rates, while I'm as dependant as anyone else on using Gmail to help save time with its quick replies to messages.

That, to me, sums up the future for AI in the next few years—it'll be something that gradually creeps into the phones, search engines, and operating systems we use, helping us out in incremental ways, until we don't realize we're dependent on it, and couldn't live without it.

In conclusion, I don't think AI is quite the same as snake oil, but I do think that it's an evolution of how we operate online, rather than the massive upheaval that the industry seems to want to sell us on. Your feature did a good job of explaining this, although the long-term implications of AI remain scarily apocalyptic, especially combined with the power of quantum computing. Just let me finish *Alan Wake 2* before a sentient computer starts unleashing *WarGames* on us, okay?

–P. Lewis

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS:
Thanks for your comments on our AI issue, and I understand some of your apprehension about AI in general! The feature also included a list of useful AI applications that are available today, and I was amazed at how many services there are, many of which are charging $10 a month or more to use. I imagine that many of these will be bought or copied by the likes of Google, Apple, or Microsoft, and rolled up into one service we pay a monthly fee for, or perhaps get for a year when we buy a new laptop or phone.

Like you, I am noticing more and more AI features just becoming the default—this week it was Google rolling Gemini answers into the top of the search results on relevant queries and questions. We're in the AI era now, whether we like it or not.

THE BUILDS

THIS MONTH'S STREET PRICES...

BUDGET

AH, 2024, a time of change. Here we have it folks, a new generation of CPUs. The cycle continues, and AMD has finally debuted its latest Ryzen 9000 generation chips to the masses. Only one problem: they're a bit underwhelming on the price to performance front, and they're currently beleaguered with some pretty significant Windows bugs, too, causing performance to be sub-optimal. This may change, but for the time being, given how great value the 7000 series is, we're sticking with the last gen.

This does mean that prices for everything else are ping-ponging around faster than a bat out of hell, which makes it particularly difficult to price these builds up. Nonetheless, we've come up with a list of six fine builds.

Our budget systems have remained relatively solid at this price point, though we have had to make a few changes. Our AMD build received a new motherboard, as the price had bounced up considerably. Alongside that, we've also picked up a new set of memory, as the Patriot Viper kit we were using had shot up by close to $30. Instead, for $4 more than last issue, we've grabbed Silicon Power's XPower Pulse gaming sku, which comes in at $4 more than last month, but beats last month's kit, boosting speeds to 6,400 MT/s and dropping the CAS latency from 36 to 32.

For our Intel system, things have stayed relatively stable, too. Our RX 7600 had increased by around $20 or so, so we've gone for the ASRock Challenger OC variant instead to offset that. In better news, the 12600KF is also down in price, and we've paired that with NZXT's T120 RGB air cooler for a nice little saving.

One other thing to note: the Corsair 4000D is back up to its $90 price mark. Given how old this is getting at this point, and pricey it remains, this should hopefully be one of the last months we see the aging giant before swapping it out for something a little more modern.

AMD INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$90
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$62
Mobo	MSI PRO B650-S WiFi NEW	$140
CPU	AMD Ryzen 5 7600	$182
GPU	ASRock Challenger D Arc A750 8GB	$190
RAM	32GB (2x16GB) Silicon Power XPOWER Pulse Gaming @ 6400 C32 NEW	$93
SSD 1	512GB ADATA Legend 840 PCIe 4.0 M.2	$45
SSD 2	1TB Kingston NV2 PCIe 4.0 M.2	$56
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$890**

INTEL INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$90
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$62
Mobo	ASRock B760M-HDV/M.2 Micro-ATX	$98
CPU	Intel Core i5-12600KF* + NZXT T120 RGB NEW	$196*
GPU	ASRock Challenger OC RX 7600 8GB NEW	$250
RAM	32GB (2x16GB) Silicon Power Value Gaming @ 6000 C30	$89
SSD 1	512GB ADATA Legend 840 PCIe 4.0 M.2	$45
SSD 2	1TB Kingston NV2 PCIe 4.0 M.2	$56
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$918***

IN THE MID-RANGE DEPARTMENT, price fluctuations have hit us hard. In our AMD build, the case shot up by $40, the motherboard by $10, the cooler by $10, and the GPU by $5. Even so, there are big wins here, too, mostly spurred on by the launch of the new 9000 series. The Ryzen 7 7700X fell in price by $40 at the time of writing, and our memory kit also fell by $12 as well—a nice little win.

Word of warning: there are a lot of NZXT H7 Flows out there at this point, all with slightly different prices. The best thing you can do is go onto PC Part Picker, and find the flow that's right for you, certainly on the pricing front. Some shops are actually selling the H7 Flow RGB for just $100, while others are selling the standard model for $120 as stock wanes. It's likely that this is just a case of the 2022 and 23 models becoming EOL, and the 2024 variant taking its position here instead. Regardless of which one you pick, you shouldn't have any compatibility issues with the build below.

For our AMD system, we also dropped the cooler down to the U12S Chromax Black in order to save a little extra cash here, too, as our $100 pick from last month had been pushed up to $110. For Intel, changes are similarly minimal. We picked up the Zotac Twin Edge OC 4070, as it's the cheapest 4070 available to date, and changed out the AIO as well just to save a touch of cash, as the Cooler Master had also seen a price increase here, too.

Interestingly, Intel's 12700K is actually holding steady at that $230 mark, making it an incredibly tantalizing proposition, despite being a few years out of date. We're still not ready to recommend 14th gen again just yet, at least until we've tested the latest BIOS micro-code updates that are launching this week.

AMD INGREDIENTS

PART			PRICE
Case	NZXT H7 Flow		$100
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold		$95
Mobo	Asus Prime X670-P WiFi	NEW	$210
CPU	AMD Ryzen 7 7700X		$219
Cooler	Noctua U12S chromax.black	NEW	$90
GPU	Asus Dual OC Radeon RX 7700 XT 12GB		$395
RAM	32GB (2x16GB) Silicon Power Xpower Zenith Gaming @ 6000 C30		$88
SSD 1	1TB Lexar NM790 w/Heatsink M.2 PCIe 4.0		$80
SSD 2	2TB Silicon Power UD90 M.2 PCIe 4.0 SSD		$103
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)		$32

Approximate Price: $1,412

INTEL INGREDIENTS

PART			PRICE
Case	NZXT H7 Flow		$100
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold		$95
Mobo	MSI Z790-S Wifi ATX		$159
CPU	Intel Core i7-12700K		$230
Cooler	Enermax Liqmax III 360 ARGB - 360mm AIO	NEW	$80
GPU	Zotac Twin Edge OC 4070 12GB	NEW	$530
RAM	32GB (2x16GB) Silicon Power Xpower Zenith Gaming @ 6000 C30		$88
SSD 1	1TB Lexar NM790 w/Heatsink M.2 PCIe 4.0		$80
SSD 2	2TB Silicon Power UD90 M.2 PCIe 4.0 SSD		$104
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)		$32

Approximate Price: $1,498

UNSURPRISINGLY, we've seen the biggest price spikes on our Turbo systems this month, particularly with our AMD build. Both the motherboard and the PSU shot up by a staggering amount. The PSU alone went from $125 to $190, and the motherboard from $280 to $311. Memory similarly increased by around $20 as well. You just can't win.

Still, there are some big wins here, too. Again with 9000 launching, the 7000 series has taken a tumble, and the price of the 7950X has fallen by nearly $100 compared to last issue, making it an absolute steal by comparison, particularly given the 16 cores and 32 threads you get for that investment. Even if you want to upgrade later down the line in a year or so, the X670E AM5 boards still have compatible with all the 9000 series chips, so you'll be in more than a stable enough position for that.

Similarly, RAM prices have shuffled around a bit here, too, and we weren't huge fans of running a 6,000 MT/s kit on the top-spec chip, so instead we've got for an upgrade and paid $5 on last month, picking up the Patriot Viper Venom 6400. It's got tighter CAS latencies as well, dropping those real world speeds.

Then we come to our Intel build, and that was a whole lot of fun. PSU? Changed. Motherboard? Changed. Cooler? Changed. CPU? Changed. Memory? Changed. Yep, the list went on and on and on, mostly as we saw big price hits across the range. The CPU alone went up by $80 this month, so we dropped down from the KS to the KF instead. You do lose the iGPU by doing that, but the price drop is very big, and saves us $20 on last month—this would have gone up to $100 if we'd have kept the KS.

Similarly, our cooler pick, the Corsair H150i Elite, had shot up by $50 as well. The company's going through a bit of a stock market downturn right now, so we suspect that this may be a result of that. Still, we managed to keep it fairly tight, and you still get that 360mm AIO to keep things chilled.

AMD INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$145
PSU	Super Flower Leadex VI Platinum Pro 1000W - 80+ Platinum NEW	$130
Mobo	Gigabyte X670E Aorus Pro X - AM5 NEW	$280
CPU	AMD Ryzen 9 7950X	$449
Cooler	NZXT Kraken 360 - 360mm AIO	$180
GPU	ASRock Phantom Gaming OC Radeon RX 7900 XTX 24GB	$909
RAM	64GB (2x32GB) Patriot Viper Venom @ 6400 C32 NEW	$180
SSD 1	2TB MSI Spatium M570 HS PCIe 5.0 M.2	$200
SSD 2	2TB Lexar NM790 PCIe 4.0 M.2	$135
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $2,640

INTEL INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$145
PSU	Super Flower Leadex VI Platinum Pro 1000W - 80+ Platinum NEW	$130
Mobo	Gigabyte Z790 Pro X ATX NEW	$280
CPU	Intel Core i9-12900KF NEW	$280
Cooler	Asus ROG Strix LC II ARGB 360mm AIO NEW	$133
GPU	PNY Verto Overclocked RTX 4080 Super 16GB	$960
RAM	64GB (2x32GB) Patriot Viper Venom @ 6400 C32 NEW	$180
SSD 1	2TB MSI Spatium M570 HS PCIe 5.0 M.2	$200
SSD 2	2TB Lexar NM790 PCIe 4.0 M.2	$135
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $2,475

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 *Mac|Life*, 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