SURFACE 7 TEARDOWN PG. 58

CHATGPT vs. **GEMINI** PG. 92

BEST **MINI MOTHERBOARDS** PG. 12

MAXIMUM PC

MINIMUM BS • NOVEMBER 2024 • www.maximumpc.com

RYZEN 9000 UNLEASHED!

Our full guide to AMD Zen 5

- Ryzen 9 9900X built & benched
- New X870 motherboard tested
- **PLUS!** The fastest SSD ever

STEP-BY-STEP GUIDE BUILD THIS PC! PG. 16

NEW INTEL & QUALCOMM LAPTOPS PG. 46

THE ETHICS OF AI PG. 32

FUTURE
VOL 29, NO 11

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

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

EFF.ORG

ELECTRONIC FRONTIER FOUNDATION

Protecting Rights and Promoting Freedom on the Electronic Frontier

inside

NOVEMBER 2024

SCAN TO GET THE TOM'S HARDWARE WEEKLY NEWSLETTER

QUICKSTART

16 RYZEN 9000 UNLEASHED

© MSI, 11 BIT STUDIOS

IN THE LAB

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: Hope Corrigan, Richard Devine, Nate Drake, Andy Edser, Nick Evanson, Ian Evenden, Jeremy Laird, Christopher Livingston, Chris Lloyd, Zak Storey, Nick Peers, Jarred Walton
Production Editors: Katharine Davies, Steve Wright
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Photography: Neil Godwin, Olly Curtis, Phil Barker
Cover Photo Credits: Google, OpenAI, MSI, ASRock, 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
Brand Director: Faisal Alani
Content Director: Marc Chacksfield
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 November 5, 2024

EXTRA DIGITAL FEATURES

AUDIO FILE | PHOTO GALLERY | VIDEO FILE

Guy Cocker

RYZEN TO THE CHALLENGE?

FOR THE LAST FEW YEARS, it's felt as though CPUs have been the bridesmaid, never the bride. With the rise of AI, cryptocurrencies, and gaming during and after the lockdowns, it's felt like all the love has been reserved for GPUs. There have been so many stories to cover—the rising price of premium graphics cards (the RTX 3090 launching at $1,499, then the 4090 for $1,599). Intel entering the market with its Arc series. AMD achieving dominance in the nascent handheld space started by the Steam Deck. The list goes on.

But this year has become all about the CPU. Qualcomm shook things up with the first Arm-based chips to run Windows—not just that, but powering the next generation of Microsoft's OS in the form of Copilot+ PCs. And now Intel is fighting back with its Core Ultra Processor Series 2, which Intel says has even better performance and battery life than its Arm-based competitor, but on the x86 platform, meaning unrivalled software support. If its promises (which you can see in our lead news item on page 8) prove true, it's made Qualcomm's chips if not obsolete, then way less desirable.

Which brings me to this issue's cover build, focusing on the latest generation of AMD's Ryzen desktop processor, the 9000-series. Team Red's new CPU launch doesn't offer the same sort of generational leaps that we're currently seeing in the mobile space, but Zen 5 is still an important launch for AMD. First of all, it's an easy upgrade for owners of 8000 and 7000 chips, as the new chips work in previous-generation motherboards. Secondly, while the performance leap is relatively minor, they are much more efficient than previous Ryzen chips, and a lot more efficient than Intel's power-guzzling 14th and 13th-gen chips.

But is it worth upgrading to these new CPUs, as well as the new X870 motherboards that arrive with them? That's the question Zak Storey sets out to answer with his build. His new suite of benchmarks also cover how strong a chip the Ryzen 9 9900X is for gaming, rendering, and AI tasks, as well as how it needs to be cooled and how much power it consumes. The huge build feature starts on page 16, and be sure to watch our video of us discussing the build on page 17.

This month, I went to Berlin, Germany, for the IFA 2024 show, which had an abundance of next-gen laptops that we can't wait to get our hands on. I report on the best on page 46. Zak also examines the ethics of AI, and Nate Drake looks at the history of something we pretty much take for granted every time we use our computer: encryption.

We also have a stacked tutorial section, starting with a look at Microsoft's Qualcomm-powered new laptop and tablet in conjunction with iFixit. Thankfully, they've come on leaps and bounds in terms of repairability. Nick Peers also shows you how to find anything using the amazing Flow Launcher, rip the perfect media file to consume on whatever device you want, and write your first (or next) novel.

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

The Eagle has landed

Intel's Lunar Lake arrives on the mobile market

INTEL'S renaissance continued with the launch of Lunar Lake at this year's Innovation For All conference in Berlin (the biggest consumer electronics show in the world, and we were there—see our report, pg. 46). Lunar Lake brings together many of the technologies Intel has been working on over the last few years. We have all-new microarchitectures for both types of cores, and a new GPU using Battlemage (Xe2). There's a new NPU, embedded memory (a first for Intel), new 3nm TSMC tiles, and more. This is a new chip from head to toe. Every processor launch is hailed as important, but this one has a lot to prove. Billed as the complete package, Intel boldly claims "we nailed it."

At the heart of Lunar Lake are four Lion Cove P-cores and four Skymont E-cores. This means Intel has pitched the chip at thin and light laptops with fairly modest prices, but which also sell in big numbers. Lion Cove offers an estimated IPC bump of 14 percent, which is decent for a generational jump. The Skymont cores offer an estimated 68 percent jump in floating point performance. The graphics are a claimed 50 percent faster than Meteor Lake, and capable of running AI at 67TOPS. The NPU is fourth-gen and capable of 48TOPS.

A group of very pleased Intel people at the Lunar Lake launch event, who were clearly proud of their work.

Intel is particularly pleased with power consumption, which is claimed to be half that of Meteor Lake. The average maximum draw is 37W, with a base of 17W, and a minimum of 8W. Intel rates Lunar Lake as having 1.2 times the performance per watt over the Snapdragon X Elite, and 2.29 times the performance against Meteor Lake. As a result, it's the most power-efficient x86 processor ever. Laptops tested by Intel managed over 20 hours of active workload.

There are nine versions in all. The top model is the Core Ultra 9 288V, which is the only one that steps outside the impressively abstemious power envelope. It has maximum speeds of 5.1GHz for the P-cores, and 3.7GHz for the E-cores. They all have either an Arc 170V or 130V GPU, 12MB or 8MB of Smart Cache, and 32GB or 16GB maximum memory. The fixed core count and restrictive power envelope means there is little room for variation. The split is probably as much a result of the binning process as it is the marketing strategy.

Intel appears to be very happy with Lunar Lake's performance. It presented bar charts showing it beating its rivals running Geekbench and Cinebench 2024 by over 20 percent. It also claims the Battlemage GPU bests AMD's Ryzen AI 9 HX 370 by 16%, and is the fastest integrated GPU in the industry.

Meanwhile, on Intel's plans for its process nodes, we may have hit a snag. It promised five nodes in four years, and the next due is 20A, which moves from FinFET to RibbonFET (a gate-all-around design, a major shift in transistor gate construction). After that, 18A is due, a refinement of 20A. The process nodes are developed in parallel, and it seems Intel was so pleased with the progress of 18A, it's decided to skip 20A and shift resources to 18A, which is on track for next year. This brings useful cash savings, but leaves Arrow Lake, the planned recipient of 20A, without a foundry. TSMC, which makes Lunar Lake's processor tiles, will get the business. Instead of 20A Intel silicon, Arrow Lake, Lunar Lake's desktop cousin, will get 3nm ones from Intel's rival.

18A will put Intel back at the edge of chip tech development, and it's betting the farm on it. Intel's foundry business lost $7 billion last year, and only hopes to reach break even in 2027. Reportedly after extensive tests on 18A wafers, Broadcom (a potential customer) wasn't happy, concluding they aren't ready for large-scale production. Intel's more optimistic, saying 18A was "powered on, healthy, and yielding well, and we remain fully on track to begin high volume manufacturing next year". Broadcom then publicly denied it made such an evaluation.

There's no denying what Lunar Lake has delivered, however. Chips should be available by the time you read this, and there are 95 new notebook designs from 25 top manufacturers due soon. **–CL**

> "Laptops tested by Intel managed over 20 hours of active workload

HERE COMES THE 5090

NVIDIA IS PREPARING its next-generation graphics cards, the Blackwell-based 50-series. Frankly, we were expecting much more at this stage. Nvidia has denied rumors of delays, and says it's made changes to the photomask (the chip pattern used in lithography) to increase yield. This does sound like a delay to us. Blackwell's main job is as a massive AI accelerator in racks, which is where Nvidia will make the billions of dollars the market expects to see in its results. Like it or not, consumer graphics cards come second.

Unconfirmed lists show that there will be five desktop GPU variants, from the GB200 to 207, covering the GeForce RTX 5090 down to 5060. Hopes are high, with fat data pipes, high clocks, and huge jumps in rasterization performance of 50 percent or more.

There are conflicting rumors on power consumption, initially thought to go to 600W for the 5090 Ti, which is the limit for the 16-pin 12VHPWR connector. However, 550W is now favorite. That's still hot, and the whole range will be power hungry. We'll probably have to wait until January's CES before we see the first desktop cards, the 5090 and 5080. **–CL**

INTEL AND AMD AI PCs THIS FALL

Red and blue teams set to join AI party

MICROSOFT WROTE the specification for the AI PC, which must have made for frustrating reading at Intel and AMD. It required an NPU with a performance of 40TOPS, and Microsoft exclusively used Qualcomm's Snapdragon X Elite chips on the first systems. This must have been galling for Intel, as that arbitrary number left Intel's Meteor Lake, which can only muster 34TOPS, out of the mix. It almost looks deliberate. This fall, the blue and red teams can get into the game when AMD's Ryzen AI 300 chips and Intel's new Core Ultra 200V series get to join the party and we'll get a slew of new CoPilot+ systems.

Meanwhile, Qualcomm has announced the Snapdragon X Plus, a slightly cut-down version of the Elite running eight-cores. Still powerful, but usefully more affordable, it retains the Elite's 45TOPS AI power. Qualcomm looks to have effectively cut itself a nice slice of the mobile market.

The focus on the NPU in these chips is disingenuous. AI functions can run on any CPU, and if you have a powerful GPU, you also have a powerful AI accelerator which will beat current integrated NPUs on performance. Microsoft will offer GPU support in a future Windows update. **–CL**

Tech Triumphs and Tragedies

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

TRIUMPHS

BLUETOOTH 6.0 FINDS IT
The first big update in Bluetooth since 2016 brings device location to "centimeter-level accuracy over considerable distances."

ZETTASCALE COMPUTER
Japan plans a supercomputer capable of 1,000 exaFLOPs, 1,000 times more powerful than Oakridge's Frontier.

BIOHYBRID ROBOT
Researchers at Cornell have built a bot that uses a fungus as a sensor.

TRAGEDIES

SONY'S AAA BOMB
Hero shooter *Concord* sold so badly it was withdrawn after 11 days, and refunds issued. Estimated loss: $200 million.

1990S HDDS BRICK
Hard drives fail if they aren't used, with music companies finding a 20 percent fail rate when trying to access old recordings for remastering.

APPLE TO DITCH USB-A
The new Mac Mini is due. Word is that it won't sport USB-A.

AMD'S KRAKEN POINT

New Zen 5 mobile aims at Lunar Lake

AMD'S RYZEN AI 300 series chip appeared this summer to ride the wave of AI PCs, although Microsoft has only just got round to adding them to its CoPilot+ initiative. AMD started at the high-end, and the resulting laptops are expensive $1,000-plus affairs. To get at the bulk of the market, AMD needs something more affordable. Enter Kraken Point, an eight-core chip, with four full-fat Zen 5 cores, and four smaller Zen 5c ones. This means systems should dip well into three figures, where the mass market is. It'll also be a competitor for Intel's Lunar Lake and a candidate for handhelds (Steam Deck 2?).

AMD also has a pair of new tricks for its 300-series to boost gaming. Variable Graphics Memory enables you to allocate up to 75 percent of the memory to your integrated GPU. You'll need at least 32GB to start with, as you still need to leave 16GB for the CPU. The results are variable—not everything will see a benefit, and some games will suffer. However, it will allow you to play titles that otherwise wouldn't run at all due to memory requirements. The second new feature for the mobile chips is frame generation, or Fluid Motion Frames 2, to be exact, which makes the jump from discrete graphics cards. This needs 50fps to work effectively, which can be a tough ask for big titles on many laptops. Both the new features are available on AMD's Adrenaline software now.

There's more AMD mobile silicon to come next year, too, with Strix Halo, a proper high-end chip with 16 full Zen 5 cores and a 40-CU GPU. Fire Range is also due, again with 16 cores, but a much less capable GPU, and is being aimed at productivity tasks rather than gaming. The competition in the mobile market, which is often very profitable, is getting fierce. Current predictions are that nearly all the growth in the PC market over the next few years will be in the laptop sector. **–CL**

PLAYSTATION 5 GOES PRO

Play everything in fidelity mode at full speed

WE WON'T SEE a PlayStation 6 for a couple of years, but to tide us over, Sony has the PS5 Pro. It's pretty much a repeat of the PS4 Pro: a mid-term bump in power, in this case, a 45 percent increase in rendering speed. The new GPU has 67 percent more Compute Units and 28 percent faster memory. According to Sony, the PS5 Pro graphics system is built on three pillars. The first is that upgraded GPU. Second is the advanced ray tracing said to be two to three times faster. Lastly, there's PlayStation Spectral Super Resolution, a fancy name for AI-powered upscaling. This helps the PS5 Pro reach 60fps on many games that would only hit 30fps in fidelity mode on the PS5. We also get Wi-Fi 7, and double the storage at 2TB. There's no mention of the CPU, so it's more than likely the same, possibly with a slight clock tweak. The Pro only comes as a digital edition. If you need an optical drive, that's a separate purchase at $79. The case gains a set of black stripes across its midriff, otherwise it looks the same as the PS5 Digital Edition.

The upshot is now you don't have to choose between performance and fidelity modes. Games may look better in fidelity, but according to Sony, three quarters of PS5 players prefer performance mode. The new model is yours for $699. This is serious money for a console, and $250 more than a standard PS5. The PS4 Pro was only $100 more, and that was after the original's launch price had dropped by the same amount. The PS5 Pro doesn't look good value if you already have a PS5. Even attracting newcomers may be tough. This launch needs a bundled game or two optimized to make the most of the new hardware, or even a new game. Hardcore fans will love it though—Sony is dominating the console war. **–CL**

For $250 more, you don't have to compromise on frame rates or visuals.

Radeon quits high-end cards

JANUARY'S CES should see AMD launch new Radeon cards sporting the RDNA 4 architecture. These are destined to go against Nvidia's Blackwell 50-series, which could be tough. AMD is a distant second to Nvidia on desktop cards, taking about 17 percent of the market last year. Nvidia makes the fastest cards, and has done for some time. There've been rumors of AMD changing tack, aiming at the mainstream market, and leaving the low-volume, high-end market to Nvidia. In a recent interview, AMD's Jack Huynh said: "We want to build gaming systems for millions of users." AMD is going to concentrate on what it does best: good value budget and mid-market cards, with the ambitious target of a 40 percent market share. The hope is that developers will start optimizing for AMD cards.

This month also saw AMD's launch of the Ryzen RX 7800M, a discrete mobile GPU, and the second to feature RDNA 3, this time a Navi 32 chip. It has 60 compute units and 3,840 stream processors. The base clock is 2,145MHz. There's 12GB of memory on a 192-bit bus, with a top transfer rate of 432GB/s. It runs about as fast as the similar RX 7700 XT, although the memory bus has been scaled back from 256-bit. It's said to be more than a match for an Nvidia RTX 4070 laptop GPU. The 7800M specs are also suspiciously close to the chip inside the new PS5 Pro. **–CL**

AMD decides to focus on market share, for now.

Window 11 is gaming favorite

Win 11 has become the most popular gaming platform on Steam. The August numbers have Win 11 on 49.17 percent of systems, pipping Win 10 at 47.09 percent. Every new PC has Win 11, so Win 10 was always going to be overtaken, but it's remarkable that it's taken three years. Especially as Win 11 is a free upgrade—clearly not a hugely popular one, not helped by the hardware requirements. Win 10 suffers from the 'good enough' problem that afflicted Win XP, which was supposed to die in 2009, but its popularity meant support continued until 2019. Officially, Win 10 has about a year to go before its EOL date. Microsoft will hopefully offer an extension. **–CL**

Most AI fails

The tech industry is falling over itself in the rush to get AI into everything, and not always successfully. About 80 percent of AI projects fail, according to non-profit research foundation RAND. This is twice the rate of non-tech start-ups, and a third higher than the already disappointing tech start-up rate. There are five root causes, which aren't unique to AI, including misunderstanding the problem, or tackling one that can't be overcome by AI. Entrepreneurs don't always understand what AI can do. Engineers also get bewitched by new tech, rather than concentrating on a practical application. The rush to AI has all the hallmarks of other such tech movements, that leave a trail of failed ventures. **–CL**

Jarred Walton

TECH TALK

US wants Nvidia & Apple to use Intel's foundries

THE US GOVERNMENT has been busy over the past year or two handing out large subsidies and grants to help boost local silicon manufacturing efforts. Intel was once the undisputed leader in the race for ever smaller and faster fabrication nodes, but after major stumbles in the transition from 14nm-class to 10nm-class technologies, TSMC passed it by.

Now Intel is playing catchup, investing billions into bringing up new nodes. Under CEO Pat Gelsinger, it has also been working to take on TSMC directly by offering its foundry services to other companies. Things haven't exactly gone smoothly.

During a meeting with the US Department of Commerce Secretary Gina Raimondo, Gelsinger expressed frustration with US companies' reliance on TSMC. Specifically, he's talking about Apple and Nvidia—though this also applies to AMD and other smaller companies. After the meeting, Raimondo has apparently been holding meetings with some of these companies, encouraging them to use US foundries to produce advanced AI and other chips.

There's ample reason for the US to want to be involved. First, Taiwan is in a sketchy geographical location, next to China, with the powers that be often mentioning how it should rejoin the rest of the country. If that happens, imagine the chaos that could result for the top AI processor manufacturers that almost exclusively use TSMC. TSMC now has a plant being constructed in Arizona, which could help, but diversification remains a key interest.

Imagine if TSMC doesn't execute as well as it would like on a future process node. Perhaps Intel Foundry Services (IFS) will regain the lead. If no one is actively using IFS, though, it would become less likely to compete, and even if it does surpass TSMC, it could take years for products that are already in the pipeline to make the switch from TSMC to IFS.

> Some of the disbursement of funds has been delayed while Intel gets its act together

Intel CEO Pat Gelsinger poses with US Commerce Secretary Gina Raimondo to mark CHIPS Act funding for Intel.

The US has been investing billions of dollars into Intel and IFS via the CHIPS Act, and it would like that money to not be wasted. Some of the disbursement of funds has been delayed while Intel gets its act together—it recently let go of 15,000 employees in an effort to cut costs. But while CHIPS Act funds are being spread around, it's clear that the predominantly US-based Intel is one of the government's key strategic partners for processor manufacturing.

Watching how things play out with Intel, IFS, and future planned nodes will be interesting, as Intel is still playing catchup. However, Nvidia CEO Jensen Huang said that the GPU maker could shift its fab if needed. "In the event that we have to shift from one fab to another, we have the ability to do it. We won't be able to get the same level of performance or cost, but we will be able to provide the supply." It may be true that at present IFS can't match TSMC on performance or cost, but partnering with big companies like Nvidia and Apple would help improve its standing.

Intel plans to split off IFS into an independent subsidiary. That will likely be key in procuring business from competitors such as Apple and Nvidia. The new IFS will have a separation of profit and financial reporting, with its own independent board. It will report to Intel's CEO, but will offer better opportunities to explore external funding.

There's been a lot going on at Intel for the past couple of years, and it's been a rough ride at times. However, many of the issues have likely been years in the making and can't be blamed on Gelsinger, who's only been at the helm since February 2021. He's made a lot of waves and shaken things up, but if it results in a leaner, hungrier Intel, it could end up being successful.

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

THE LIST

THE BEST MINI-ITX MOTHERBOARDS

FANS OF SMALL-FORM-FACTOR PCs are a dedicated bunch, and a Mini-ITX motherboard is the absolute superstar of such builds. Yes, you may want to try to fit a PC into a case not much larger than a shoebox, but that doesn't necessarily mean you want to compromise on performance, I/O, or upgradeability. These are our top picks for the best options for both Intel and AMD builds, and at a variety of budgets.

3 ASUS ROG STRIX Z790-I GAMING WIFI

Our number 3 pick is the Intel version of our number 1 mobo—it's only lower because LGA 1700 is nearing the end of its life, lacking the upgrade path of its AMD sibling. For now, though, if you want a board that can accept the most powerful 12th, 13th and 14th-gen CPUs, the Z790-I Gaming WiFi is an excellent choice.
$470, www.asus.com

5 MSI MPG B650I EDGE WIFI

This is a classy midrange option that's perfect for anyone who wants a feature-rich Mini-ITX motherboard without spending the big bucks on a more expensive X670 board. It's still more than up to running Ryzen 7000 chips and their X3D variants, as well as newer Ryzen 9000 chips, too. It also has four SATA ports, making it perfect for data hoarders.
$270, www.msi.com

2 GIGABYTE A620I AX

The Gigabyte A620I AX is an excellent option for those on really tight budgets. That doesn't mean it skimps on specs, though, with the inclusion of Wi-Fi 6E and a 2.5G LAN port on the networking side particularly noteworthy. Pair it with a Ryzen 8000-series CPU and you have the most affordable path to 1080p gaming you can get right now.
$130, www.gigabyte.com

4 ASROCK B760I LIGHTNING WIFI

This is a B760 board that could easily substitute for a Z790 one. It can handle a top-end Intel 14900K CPU and DDR5-8000+ memory out of the box, no BIOS updates necessary. It only misses CPU overclocking and PCIe bifurcation support, but you would need to spend significantly more to get those from a Z790 board.
$185, www.asrock.com

1 ASUS ROG STRIX X670E-I GAMING WIFI

This is the best Mini-ITX AMD board by far. With Gen 5 GPU and SSD support, plus USB4 and good-quality audio, the X670E-I is a good step up from nearly every other Mini-ITX board. It's also got a VRM capable of powering the high-core-count Ryzen 7000 and new 9000 series CPUs, perfect for anyone eyeing up Zen 5.
$440, www.asus.com

Jeremy Laird

TRADE CHAT

PS5 Pro makes the PC more relevant than ever

THE PS5 PRO has arrived. It's far, far more powerful than the plain old PS5, but it's also a lot more expensive at $700. Does the PS5 Pro make gaming PCs look like good value again? Or are GPU prices still so high, you'd have to be very serious about PC gaming to overlook Sony's latest console?

Boil it down to basics, and the PS5 Pro is a $200 premium for a GPU upgrade. The rest of the PS5 Pro's hardware hasn't changed much. But the uplift from $500 to $700 puts the Pro into territory where it might just have to fight to justify its existence. As a gaming platform, a $500 PS5 versus a $500 'gaming' PC was always a bit of a non-starter. It's just not enough money to build a decent entry-level rig. But $700? Well, that's a little more plausible.

In graphics rendering terms, you'll struggle to match a PS5 Pro with a $700 PC. The new PS5 Pro has a revised GPU with 60 AMD compute units. That's the same as a desktop AMD Radeon RX 7800 XT, which goes for $475 and upward on its own, and to which you need to add a case, motherboard, CPU, memory, storage, and power supply, at minimum.

Of course, 60 PS5 Pro compute units aren't directly comparable to 60 desktop AMD compute units. The desktop GPU will almost certainly have the higher power budget and clock speed, probably more cache memory, and so on.

On the other hand, the new GPU in the PS5 Pro has upgraded ray-tracing hardware that may well be more powerful and efficient than the RDNA 3 hardware in the RX 7800 XT. And while the desktop GPU will likely have that bigger power budget, it's always a case of diminishing returns as you approach the top end of the power scale.

You also have to factor in the vagaries of comparative CPU performance, what with the PS5 Pro's elderly AMD Zen 2 cores. It has the same eight Zen 2-spec CPU cores as the original PS5 at the same 3.5GHz.

> Arguably, even an old Zen 2 CPU is fine for most games. But most games are not all games

Sony's new console brings a major GPU upgrade, but it actually makes PC gaming look more appealing.

Over in the PC, we're three generations ahead on Zen 5, and with clock speeds well over 5GHz. Of course, you're not going to build a brand new PC with a top-spec Ryzen 9000 chip with Zen 5 cores, plus a good GPU, and all the other bits for $700. But you might well be building one with a Ryzen 7000 Zen 4 chip, or an Intel 12th or 13th Gen CPU, all of which will blow a hole in that Zen 2 CPU in the PS5 Pro.

Arguably, even an old Zen 2 CPU is fine for most games. But most games are not all games, and even a mid-range modern PC processor has the measure of the PS5 Pro.

Anyway, the PS5 Pro's potential frame rate performance will likely fall somewhere between the RX 7800 XT and RX 7700 XT in terms of relative PC GPUs. That looks like game over for the PC. After all, you ain't buying a pre-built gaming rig with an RX 7800 XT for $700.

But here's the thing. The PC isn't just available as a pre-built rig. The PS5 Pro also has to take on existing PCs with a GPU upgrade. Or maybe a GPU and CPU upgrade. Heck, maybe even a GPU, CPU, motherboard, and RAM upgrade.

Let's take the toughest of those challenges, the full CPU, GPU, motherboard, and RAM upgrade. A quick spin on PCPartPicker nets an Intel i5-12600KF CPU, RX 7800 XT graphics, ASRock Z690 Pro RS ATX board, and 32GB of Corsair RAM for just over $800.

Yup, that's more than the PS5 Pro, but for my money, it's not even a close contest. Because then you have an actual PC, that wondrous and endlessly configurable and ever-upgradable machine, with an unbeatable back catalog of games and broader capabilities that go so much further than gaming. So, the PS5 Pro is a very nice upgrade. And it's still great value despite the price bump. But it makes PC gaming more relevant than ever.

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

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> Tracker blocking
> Best free VPN
> Remote iPad control

Email alternatives

I was advised to use DuckDuckGo Email Protection (https://duckduckgo.com/email/) to improve my email privacy, so set up free Duck.com email addresses for myself and my family. Sadly, they stopped working. DuckDuckGo told me my email provider has started rejecting the emails as spam. Is there any way to overcome this, as we've given our Duck.com address to many companies? Or is there an alternative?

–Bob White

Free VPNs provide good—if limited—online privacy and protection.

THE DOCTOR RESPONDS:

How ironic that your email provider has decreed that messages sent from the DuckDuckGo servers are spam, given its sole job is to strip trackers and other spam-like behavior from emails. The Doc assumes that by 'rejected', your provider isn't forwarding the mails (while marking them as spam) but bouncing them back to the sender. The first option is a simple fix: Check your spam or junk folder and look for options to mark the emails as safe or trusted.

This leaves two options: Choose an alternative email forwarding provider, such as Addy.io (https://addy.io), or reroute your Duck.com emails to a different email provider. The first option probably isn't practical as the likes of Addy.io exist solely to hide your email address from third parties, not vet the emails to remove any tracking information.

That leaves you with the option of forwarding your Duck.com emails to a different email address with a provider that won't start rejecting them. To be clear, the Doc isn't suggesting you ditch your current email address as that would be impractical. Instead, you'll use it solely for direct personal and (if applicable) professional correspondence. You'd then set up a secondary address with another provider specifically to deal with companies and other organizations. This might be a webmail provider, such as Gmail or Microsoft (Live.com), or a free mailbox with someone like Proton (https://proton.me).

Once set up, you'd either reserve a browser tab (the Doc recommends pinning this) for checking that webmail account, or use an email client like Thunderbird or Betterbird (https://betterbird.eu) that allows you to access all your email addresses from a single app.

Podman glitches

I see that Podman 5.0 has come to Debian—it was installed among other regular updates on my server. Now, after rebooting my server, I discover Nginx Proxy Manager is no longer able to communicate with one of my containers (Jellyfin) and neither can I use it to redirect subdomains to natively installed services either. What's happened?

–Robert L Campanelli

THE DOCTOR RESPONDS:

This is linked to the fact that one of the key 'breaking' changes in Podman 5.0 is its new default network stack. Rootless containers are now managed using the 'pasta' network stack as opposed to 'slirp4netns', and prior to the release of Podman 5.1, the pasta stack was unable to communicate directly with the host. This means Nginx Proxy Manager could only manage connections to containers in the same defined network. It's part of a wider change whereby Podman now defaults to its own native Netavark network stack instead of

↘ submit your questions to: doctor@maximumpc.com

the older CNI (Container Network Interface) stack.

The good news is that Debian has just updated to Podman 5.2.1 in its repositories, which fixes the problem. However, this comes with its own set of issues, namely a string of Podman services failing to start due to a database configuration mismatch error that says the "database storage graph driver does not match our storage graph driver overlay."

Fixing this is done with the following terminal commands. Debian users first need to type 'su' and enter their root administrator password to switch to root mode. Now issue these commands:

```
cd /var/lib/containers/
storage
sqlite3 db.sql
```

This should drop you into a sqlite> prompt. Now issue the following command:

```
update DBConfig set
GraphDriver = 'overlay'
where GraphDriver = "";
```

If the command runs with no error, press Ctrl-D to force quit SQLite3, then type 'exit' and press Enter to exit root user mode. You should now be able to restart all Podman services manually, or reboot, when they should now run successfully.

Fixing Nginx Proxy Manager from here is simple—if you don't need to remap ports 80, 81, or 443. Remove all port mappings from your Nginx Proxy Manager container script, then adjust the script so it's communicating on the host network and not one you created for Podman (so '--net=host'). After the container has been recreated, Nginx Proxy Manager can communicate with any container (even those within their own Podman networks) as well as any native services.

Free VPN

I like the idea of a VPN for occasional use, but I don't think I'd use it often enough to justify paying for one—or setting up one as per your WireGuard feature in the June 2024 issue. Do any VPN providers offer free plans, and if so, which would you recommend? **–Fred Byrom**

THE DOCTOR RESPONDS: Yes, a handful of VPN providers offer free plans, and three fall within your remit. The first is TunnelBear (www.tunnelbear.com), which is based out of Toronto. It offers a free tier that gives you 2GB of data per month to play with, which would comfortably cover basic needs like browsing and checking email and social media. You also have full access to most of its servers and features, enabling you to spoof your location if needed.

If 2GB isn't enough, another Canadian company, Windscribe (https://windscribe.com/features/use-for-free), offers 10GB of data, and can be used on an unlimited number of devices. On the downside, it only provides access to servers in 10 countries, including the US, Canada, and Switzerland. Those looking for complete anonymity should note it's the only VPN of the three featured that doesn't require you to provide an email address.

The last option would be if you're using the VPN solely to protect a single device when connecting to public Wi-Fi networks—including encrypted ones. ProtonVPN (https://protonvpn.com/free-vpn/server) is based out of Switzerland, offers no data limits and no ads. Its drawback is that you have no choice over which server you connect to—it's always the fastest available from a small pool based in five countries, including the US, Netherlands, and Japan.

Bookmark sync

I refuse to tie myself to a single browser, but my bookmarks are scattered across multiple browsers and devices to the point where I don't know which one has the URLs of my favorite sites. Is there a way to tie them all together that can be accessed from any browser on my mobiles or desktops? **–Michael R Atkinson**

THE DOCTOR RESPONDS: You might want to consider Floccus (https://floccus.org), a browser add-on that can sync bookmark folders (and open tabs) between Chrome (and other browsers with access to the Chrome web store, such as Vivaldi and Opera), Edge, and Firefox on any PC or Mac. There's also a free app for Android and iOS, although this doesn't synchronize bookmarks to your mobile browsers; instead you use it as a standalone bookmark manager, with the bonus of being able to tap any bookmark to open it in your phone's default browser.

Download Floccus on your primary browser and follow the prompts to set it up. Your main choice involves what medium you sync your bookmarks on: Google Drive is the simplest to set up, but set a strong password so your bookmarks are encrypted to keep them private from Google (or any hacker, should your account be compromised).

An interesting option for those who've set up their own Nextcloud instance is Nextcloud Bookmarks, which lets you sync bookmarks privately over your own hardware (albeit without encryption). There's also a Git-over-HTTPS option, but again encryption isn't supported, so you're storing them in a public space.

The final option—WebDAV share—is probably the best alternative to Google Drive. It works with any WebDAV-compatible server and supports encryption, so you could use it to sync to your web space or self-hosted server (including Nextcloud if you were looking for a more secure option).

One of Floccus's strengths is that it supports multiple profiles, so you could set these up to sync specific folders, or to restrict certain bookmarks to certain devices or browsers. It can also sync tabs between browsers. If you do decide to use it, you may need to disable your browser's own sync tool (or at least the bookmarks portion of it) to prevent potential conflicts.

Controlling an iPad

I bought my elderly mother an iPad, but she's struggling to install updates and manage the privacy settings. Is there an app that lets a trusted person look after someone's iPad? **–Alison Day**

THE DOCTOR RESPONDS: Apple's native tools for remote-controlling an iPad from another device are too restrictive to do what you'd like—basically, you can only do so from another iPhone or iPad logged into the same iCloud account and on the same Wi-Fi network. This leaves you needing to source a remote-control app that is capable of controlling iOS devices as opposed to simply allowing you to access your PC, Mac, or Android device from your iPhone or iPad, which sadly rules out RustDesk (www.rustdesk.com), which we covered in the October 2024 issue.

One tool worth exploring is Zoho Assist (www.zoho.com/assist/remote-control-iphone-ipad.html). This consists of a technician app, which you'd install on your device, and a client app that would be installed on your mother's iPad. Sadly, the free version doesn't cover iOS devices, so you'd have to pay a monthly ($12.99) or annual ($120) fee to use it. A free trial is available—given the costs involved, the Doc recommends you test the software to see whether it can meet your needs.

STEP-BY-STEP GUIDE BUILD THIS PC!
PG. 22

https://content.jwplatform.com/videos/sl7k4tVf-u2lN49He.mp4
Please type this URL into your browser if the link is broken

RYZEN 9000 UNLEASHED

Zak Storey puts AMD's latest chips and motherboards to the test

WE WERE THERE, dear reader, seven years ago, when Ryzen first landed center stage. We remember it well. A subdued press event hidden away under a hotel in San Francisco. Everything was under lock and key. NDAs were signed and lips sealed as the world's tech press shuffled into a presentation hall. Heck, we couldn't even tell folk we were getting a flight to the darn place, they were that secretive about it. Still, we all gathered at the event, then Lisa Su emerged from the shadows, and with her came a powerful CPU, one that could dominate all, and pave the way back to red-tinged processing greatness.

Ryzen's first-generation processors had landed, masterminded and designed from the ground up by Jim Keller and a crack team of silicon experts and engineers. Its very launch reimagined what a CPU could be—even should be. It was the first time we'd seen a truly modular CPU design. Chips that could be scaled up or down, intertwined with the beautifully impressive infinity fabric to deliver performance and value unlike any we'd seen prior. Intel had to scramble, and if we're honest, it's still scrambling today. No bad thing.

That was many years ago, and at long last, the 9000 series has finally graced us with its presence in a *Maximum PC* build. It's hard to believe that it's been 394 weeks since we could first even talk about its progenitor chip—394 weeks filled with ups, downs, world events, and generation after generation of processor, all of it leading up to this one moment. It's remarkable, really. In that short time, AMD's Ryzen architecture has continued to advance and develop. Transistor size has shrunk, the inclusion of iGPUs, smarter resource allocation, and intriguing new 3D cache design concepts have pushed these chips ever higher on their way to greatness. And at last, we're finally here.

So, the build? At its heart lies the AMD Ryzen 9 9900X paired with Asus's latest and greatest X870E Strix Gaming motherboard. A new chipset, a new chip, and some stellar hardware picks. We'll be honest, this is unabashedly all about the CPU this time around. It's a well-balanced build, but we're here for one thing and one thing alone: to see if AMD's latest generation is another great launch or a flash in the pan. Let's get to the parts, shall we?

SPECIFICATIONS

PART		PRICE
CPU	AMD Ryzen 9 9900X	$449
Motherboard	Asus ROG Strix X870E-E Gaming WiFi	$550
CPU cooler	360mm NZXT Kraken Elite 360	$267
RAM	32GB (2x16GB) Lexar Thor OC DDR5 @ 6,000MT/s C32	$110
SSD 1	2TB Crucial T705 M.2 PCIe 5.0	$295
SSD 2	2TB Lexar NM790 M.2 PCIe 4.0	$135
GPU	Zotac GeForce RTX 4070 Ti Gaming AMP AIRO Spider-Man Edition	$750
Case	NZXT H5 Flow RGB 2024	$120
PSU	1200W NZXT C1200 2024 80+ Gold	$200
Fans	1x NZXT F240 RGB Core Frame + 1x F120 RGB Core	$60
Total price		**$2,936**
Core price*		**$2,154**

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

PRICES CORRECT AT THE TIME OF PUBLISHING

MOTHERBOARD

Asus ROG Strix X870E-E Gaming WiFi

We know we say this a lot, but the I/O complement on this is utterly insane. But let's start with the obvious: it looks gorgeous. The array of heatsinks, pipes, blocks, and suave lighting involved in this is simply awesome. But it's more than just skin deep. The X870E-E (weird name—we won't judge), supports up to three PCIe 5.0 M.2 SSDs, each with its own massive dedicated heatsink, with the primary drive featuring a built-in heatpipe. Then there's the DIMM slots, supporting dual-channel memory up to 8,400MT/s, plus an 18+2+2 VRM setup to keep whatever 7th, 8th, or 9th series chip you're using well juiced up.

But jump around the back, and oh boy—there are eight 10GB USB Type-A ports, two 40GB USB Type-C ports (with DisplayPort connectivity), one USB-C with 10Gb/s support, one USB-C rated at 20GB with wattage watching on it, plus a 5G Ethernet port, Wi-Fi 7 as standard, not to mention mic in, line out, and SPDIF out as well. Oh, and HDMI 2.1, plus your usual clear and reset BIOS buttons.

www.asus.com

CPU

AMD Ryzen 9 9900X

Ninety dollars. That's how much extra you have to pay for AMD's latest 900X chip compared to its last-gen equivalent, although that did debut at nearly $500. Still, this is no slouch in the performance department either. It still packs in that 12-core, 24-thread configuration utilizing two CCDs and one I/O die, and also still features 64MB of total cache. Boost clock similarly remains the same at 5.6GHz max. Although that might not sound that impressive, AMD's managed to pull this off while dropping the TDP from 170W to 120W in the process.

Other items of note? Well, it's dropped from 5nm to 4nm transistors, courtesy of TSMC's N4 FinFET fab process (we've seen this already on AMD's 8000 series processors), there's been a slight bump in DDR5 memory support, with it pinging up to 5,600MT/s in dual-channel mode, and the L1 cache has been increased from 64KB to 80KB (with 48KB dedicated to data storage, rather than 32KB on the previous iteration). On the surface, there's not a lot of change. PCIe 5.0 lanes are the same, and even the iGPU has the same 2x RDNA2 CUs.

www.amd.com

STEP-BY-STEP GUIDE BUILD THIS PC! PG. 22

RAM

32GB (2x 16GB) Lexar Thor OC DDR5 @ 6,000MT/s C32

The memory pick for this system is a tricky one. We're keen to try something meaner in these builds going forward, to see what difference higher-spec memory might make, but for the time being, 6,000MT/s kits seem to be the perfect fit.

With that in mind, and because so much of the emphasis of this build is on the CPU and the motherboard, we've decided to opt for something of a low-key choice in the form of Lexar's Thor OC DDR5 kit. It actually provides one of the tightest frequency and lowest CAS latency combos out there right now, and because of that, it doesn't particularly sacrifice much in any one area. All right, it might not have the glitz and glamor punch of some of the RGB kits out there, but it'll do what we need it to do here today, and that's what's important.

www.lexar.com

CPU COOLER

NZXT Kraken Elite 360

We've given the Kraken Elite a lot of love over these last few months, and we've chosen this cooler for this build with one reason in mind: It's the same cooler we've used to test the 7900X last issue, and the 14900K in both our zero-points for this and the last issue as well. Yes, it's a gluttonous pick. It's $267 for a crystalline LCD display, plus a hefty thick rad, but it does the job, and it does it well. If you're looking for a sleek, easy, plug-and-play cooler that looks the part, NZXT's Kraken Elite series does a top-notch job. If you're after something a little cheaper, NZXT also has a few more budget-oriented options that still have an LCD screen.

www.nzxt.com

SSD 1

2TB Crucial T705 M.2 PCIe 5.0 SSD

Crucial's T705 is, without a doubt, one of the fastest SSDs right now, certainly for the PCIe 5.0 spec. In fact, it's highly unlikely we'll see SSDs run sequential storage much faster in this generation at all, as we're already approaching the theoretical maximum. We've opted for the non-heatsink variant here, ready to be plumbed straight into the X870E-E's many, many meaty heatsinks, but if you're not fortunate enough to have a ridiculously overkill motherboard like that, with its three separate PCIe 5.0 slots, each with meaty heatsink and different way to secure the drive, you can grab a heatsink version instead.

The T705 tops out sequentials in the 14GB/s range and its random 4K performance equally sits comfortably in the 300MB/s range on write speeds, making it a solid pick for an OS drive.

www.crucial.com

SSD 2

2TB Lexar NM790 M.2 PCIe 4.0 SSD

To back up our primary OS drive, we've gone with Lexar's NM790 2TB PCIe 4.0 SSD. This runs slightly less proven Chinese NAND flash at its core, but Lexar's got a solid warranty on it, and its performance numbers top out at the high end for PCIe 4.0 drives, making it a decent all-around pick for those keen to pick up a second drive for additional storage.

Why two SSD drives? Simply put, you can keep your big games and files on this drive, and your operating system on your primary drive, then re-install Windows whenever you need to for maximum system smoothness. That way, you can store all your valuable files on this secondary drive without worry.

www.lexar.com

GPU

Zotac GeForce RTX 4070 Ti Gaming AMP AIRO Spider-Man Edition

And the award for the longest graphics card name goes to.... Yeah, that's not a typo. Seriously, Zotac, what gives? Also, we know what you're thinking: RTX 4070 Ti? Not even a Super? Yup. Why's that? Simply put, the performance on this bad boy ain't half bad at all for the money, and it still very much outperforms the RTX 4070 Super (one of our favorite picks), plus the 12GB of VRAM is going to be plenty for 4K and 1440p gaming.

It's also worth noting that we're struggling to find stock on a few of the RTX 4070 Ti Super cards, and similarly, the performance bump between the two isn't really enough to justify the extra $100 or so cost in most cases, making this a far more tantalizing pick.

www.zotac.com

STEP-BY-STEP GUIDE BUILD THIS PC! PG. 22

CASE

NZXT H5 Flow RGB 2024

Last issue the H7, this issue the H5. You can read our full review of the NZXT H5 Flow RGB 2024 next month, but let's be clear: This looks slick and can pack some serious hardware in its sultry form, even being the wee little tower it is. It's got support for full 360mm AIOs (now seemingly the industry standard), plus E-ATX motherboards, full-length GPUs, and full-fat power supplies. It's even good at housing 2.5-inch drives.

And the one advantage it has over its H7 bigger brother? A stock vertical GPU bracket that fits—yep, that's right, we've gone full vertical with this little number. That is an added extra, mind you, and the case works perfectly well without it, but if you dig the look, you should know it'll set you back an additional $90 for the kit (which is a bit steep, particularly as the riser is still only Gen 4 itself; still, the quality is top-notch and in line with other NZXT products, so y'know...).

www.nzxt.com

FANS

NZXT F240 RGB Core Frame and F120 RGB Core

We've talked extensively about how swish these fans are. Simple designs, clean frames, a single cable, daisy-chainable. There's a lot to love, and installation is quick, clean, and easy.

They look stellar when illuminated, too, and are easy to control in CAM software. We've paired the frame with a standard F120 RGB Core to complete the fan complement. The case comes with its own 360 Core Frame, which we're using for the AIO; the 240 and 120 will go in the roof and rear respectively to balance the airflow.

www.nzxt.com

PSU

NZXT C1200 2024 80+ Gold

ATX 3.1: That's the key reason we picked up this little number. Plus, it gives us the opportunity to swap that 4070 Ti out at a later date for something with a lot more punch if we need to—*Warhammer Space Marine 2* isn't going to play itself. Yes, 1,200W is overkill, and you could save some cash, but 200 bucks for a PSU of this caliber is fine, particularly as it's going to live with you for a long while.

On top of that, it also comes with a 12VHPWR cable capable of delivering up to 600W, which is convenient, as we are seeing more and more rumors that this might be exactly what Nvidia is punching its new top-tier cards at.

www.nzxt.com

ALWAYS EXPERIMENT

LENGTH OF TIME: 2-3 hours
DIFFICULTY: Medium

Let's preface that difficulty rating above with a slight caveat—in a lot of ways, this is actually applicable to all manner of builds, not just this specific one—effectively, every build is only ever as complicated as you make it. You could go for an absolutely cookie-cutter, off-the-shelf system, which looks identical to the marketing material, and end up with a beautiful rig that performs well and is easy to build. However, that does lose some of its charm when you realize that the vast majority of folk using that particular case probably have it set up that way.

In our build, and given what we do here at *Maximum PC,* we have a tendency to try to push the limits of what you can do with any given chassis. In the case of the H5 Flow, what that meant is that we tried a lot of things to see what worked and what wouldn't, in an attempt to create something a little more out of the ordinary than your average H5 PC. We were the ones who made this difficult, not NZXT with its chassis design, but us (although, to be frank, it could have given us a touch more clearance in certain areas, although more on that in a moment...).

The NZXT H5 as a case is small, lightweight, and sleek. It's got plenty of airflow and all manner of perforated panels to encourage sufficient shifting of the cool stuff over your components. But it's not without fault. Clearances are tight, in no small part because of its limited form factor and size. There's an expectation on chassis manufacturers that almost all modern cases need to support a 360mm AIO as standard. In fact, we'd argue that given how toasty CPUs can get these days (here's looking at you, Intel), a 360mm AIO is almost a necessity. That is no doubt a big part of why the H5 Flow 2024 edition now features support for a 360mm, unlike its predecessor. Yet in contrast to some of the minds at Hyte, Geometric, and Fractal, for whatever reason, there's been some seriously curious design decisions made that somewhat limit the overall build itself. So, with that out of the way, let's get building.

ALL ABOUT THAT VERTICALITY

So, the chassis strip-down (**Step 1**). What's there to note here? Well, the NZXT H5 Flow, unlike the H7, requires a touch more effort to strip down to its bare bones. The window pops out in the same manner—a simple tab at the back top enables you to pop it out of place—but the rear side panel and the roof are both secured with thumbscrews, so you need to undo those to take both elements off. If you intend to install any fans in the roof, that's going to be a necessity. The front panel, fortunately, is just on poppers as well—as soon as you have that rear panel off, you can remove that, too.

A quick peek around the back reveals a fairly spartan affair in general. There's no SSD mounts on the motherboard tray, just a simple drive bay secured down under the PSU shroud with a thumbscrew. We've removed that as well for this build, just to give us a touch more cable management room. Fortunately, there's a ton of Velcro straps and cable channels in the rear as well, making it a cinch to route your cables throughout.

For this build, we decided to spice things up a bit by installing an optional vertical GPU mount (**Step 2**). Now, as we mentioned earlier, this isn't a cheap vertical bracket; it clocks in at $90 or so, making it a pricey little thing. Fortunately, it does fit in the H5 Flow (it's not compatible with the H7), and it does come with a PCIe 4.0 riser to ensure you're getting the most out of your current-gen GPU. It's also worth noting that we haven't yet saturated the PCIe 4.0 bus for GPUs just yet. There is potential that the RTX 5000 series could change that fact, but it's still too early to say. Anecdotally, from our own experience, we have experienced some coil whine from some PCIe 4.0 or 5.0 GPUs running on 3.0 or 4.0 respectively, so your mileage may vary. If that is the case, you can pick up PCIe 5.0 risers from Amazon, although they're not cheap, and you would need to replace the physical riser element in the vertical GPU bracket here.

Because we know we're going to be adding this vertical bracket, we've also taken the opportunity to strip out all of the PCIe slot covers from the back of the case as well (**Step 3**)—save these for later if you're doing it yourself—because you need to remove them first. It's also worth doing a test fit of the bracket itself at this point to ensure it will fit. Don't worry about fully securing it—that's going to be one of the last steps we do to ensure there's no conflict with motherboard clearances or similar.

ALL THE DRIVES

With our H5 Flow suitably stripped, we move on to the motherboard, and what a beauty this thing is (**Step 4**)—check out the three heatsinks at the top of the board. All for different M.2 slots, all of them massive, protruding superstructures, some complete with heatpipes and more. We're using its box as an anti-static workbench, and suitably

5

6

7

3

8

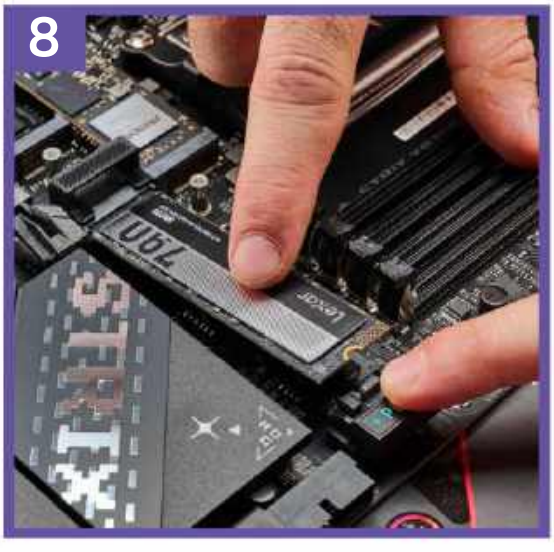

ground ourselves on something metal before operating as per usual (**Step 5**).

The topmost heatsink actually covers two of the PCIe 5.0 slots. It's got a thick heatpipe across the two—remove the four screws, pick it up, and you'll spot the twin M.2 slots underneath. The massive heatsink on the right is equally designed for PCIe 5.0, but this one has a clip-in attachment. Press the big push button and it pops right off.

We're only going to be taking advantage of two out of the three M.2 mounting solutions and slots here, but each of them is seriously impressive. The topmost slot, housing our Crucial T705, has a sliding bracket that you can move and clip into position. Pop your M.2 SSD in, making sure to align the notches, hold it down with a finger, and then push the tab into place until it clicks. Then it's a case of replacing the heatsink and its four screws, and you're done (**Steps 6/7**).

For our Lexar NM790, this is going to be going in the rightmost socket. The M.2 connector is in the center of the motherboard; we've lined and slotted that into position, and then using the push button, you simply push it down and lock it into place, before doing the same with the massive heatsink you removed earlier. The drive and sink are nice and secure. This is the quickest we've ever installed drives on to a board (**Step 8**).

With our SSDs installed, it's time to move on to the CPU. Fortunately for Zen 5 and the Ryzen 9000 series, the general CPU installation process is the same as any LGA socket. Release the retention arm, lift up the bracket, gently place the CPU into position (ensuring you line up the gold triangle with the etching on the socket, or just ensure the text is the right way around), place the bracket back down, lock the retention arm back down (**Step 9**), then the plastic cover pops off, and you're good to go.

FORGET ABOUT IT

We're almost done on the core hardware front. Next is memory, and there's no major changes here either. We're still using DDR5, and the Lexar Thor memory fits in the same as always. Lift up the tabs on the slots at either end, line up the notch, and carefully push the sticks into place (**Step 10**). It's not the prettiest of kits, but those timings and frequencies make this one of the best value kits money can buy right now.

Finally, we're on to our NZXT backplate. Find your AM4/AM5 screws and tabs. Remove one of the stock AMD CPU brackets first, either top or bottom, then place your plastic risers on top, ensuring they're oriented the correct way (one way is for AM4, the other AM5), then secure them in place with the included stand-off screws, and you're all set for later on (**Step 11**). This is one of the best things about building any AMD system. In fact, since the very early days of Ryzen, there's been little worry about included backplates for coolers; the motherboards come with a chunky one as standard, and as each side has one large plastic bracket for the stock coolers, you can remove just part of it and install half at a time. Not only does this mean you don't have to put a hand on the back of the motherboard to hold any aftermarket backplates in place, but it also means you can do this with the mobo entirely in situ in the case if you so desire.

With our prep work done, it's time to get the motherboard into the chassis. Lay the case down on its side, then carefully slide the motherboard in, ensuring it lines up with the stand-offs in the case. NZXT has a small box of accessories and screws included in the drive bay—you'll need those. Find the motherboard screws in the box and secure the board into place (**Step 12**).

With that done, we've taken the opportunity to do another test fit for the vertical GPU mount to make sure that PCIe riser hits that top slot. Again, we're not installing anything, as we've still got to install all of our bottommost connectors and cables, but this means we can avoid any unnecessary pain later (**Step 13**).

Now it's time to move on to the least fun part of the build: the cables. Fortunately, we've got plenty of room to maneuver here, and NZXT's C1200 PSU isn't exactly massive, either, certainly not for a 1200W model. We've attached the cables we need to the PSU outside of the case before sliding it into position from the rear, and fixing it in place with the screws from the accessory box (**Step 14**). You don't have much of a choice here on PSU orientation because the bottom plate is solid, meaning you need to place the PSU with the fan facing upward, or it'll choke itself to death (not good).

HERE'S WHERE THE FUN BEGINS

And this is the moment when all hell broke loose. First and foremost, if you look at the top of the case itself, you'll see there are three fan-shaped cutouts. However, NZXT says the case only supports two 140mm fans here at maximum (with low-profile memory). Because we've got problems with manufacturers telling us what we can and can't do, we opted to see whether we could secure the AIO to the roof. Now, if you have stock fans and an AIO, you could hypothetically do this, because the bulk of your fans could be secured to the slots in the top of the case. However, because we're using a Core Frame, it only has slots for four screws, one on each corner. So, to get around that, we used two cable ties and strapped the end of the Core Frame to the top end of the case. Success, we thought (**Step 15**)—360mm AIO in place, albeit a little janky, but it gives us the old "whip NZXT by saying the case can support three 120mm fans in the roof" line. Additionally, we've also had to slide the rear 120mm fan all the way


10
Lexar
AMD
RYZEN


11
AMD
RYZEN


13
NZXT
NZXT


15

down to support the radiator-fan combo here, too. Less than ideal, although not the end of the world by any length.

With that done, we moved on to installing one of NZXT's new fan controllers. This little beauty is magnetic and fits in the case like a treat. Tie all your fans together and you're on to a winner (**Step 16**). We've also gone ahead and secured the CPU cooler block to the CPU itself, cleaning off the stock thermal paste from the block, replacing the standard Intel bracket for the AMD one, and then using some of Noctua's NT-H1 thermal paste instead. Nice and easy.

It was at this moment that we decided to try to install an extra Core Frame 360 into the front to really round out the build. Sadly, the clearances were just too tight to slide it into place and secure it in position. The top radiator and the AIO would conflict with it, no matter what we did, or what orientation we ran the cables. OK, no problem, we thought. We'll swap it around, we'll run push-pull in the front, then pop the 240mm we have in the roof instead. But wait—the cutout for the front AIO and fans really only supports one set, so push-pull is out of the window. Yep, no matter what we tried, we were rebuffed every single time. In the end, the only way you can get a clean-looking build is by installing the AIO in the front, and a 240/280mm in the roof. Cookie-cutter. Why the extra fan cutout in the roof, then (**Step 17**)? Good question. We honestly can't tell you.

ALL GOOD THINGS

Thoroughly defeated, we wrapped up the rest of the cables, installing our 24-pin, EPS, and pre-routing our GPU power into position, before installing the USB cables for the fan controller, AIO, and front panel connectors. Then it was time to move on to the vertical GPU bracket.

First, click the PCIe riser into the PCIe slot at the top of the motherboard. Then slide the bracket into position, aligning it with where those PCIe slot covers were earlier. With that out of the way, you can then secure it from the rear using two of the thicker screws out of the accessory kit that came with the case. Underneath the bracket itself, you'll find there's a fairly large foot screw that you can turn to give the bracket a bit of support from the power supply cover. Twist that until the bracket is well supported.

Finally, line up your GPU and slot it into position. You can then secure it from inside the case. You might need to remove the 120mm fan to get access to the screws, but once that's done, it's simply a case of installing the 12VHPWR cable into the GPU, and you'll be good to go. Make sure the connection is solid here, too, because the last thing we want is loose sense wires sending the wrong voltage to the GPU (**Step 18**).

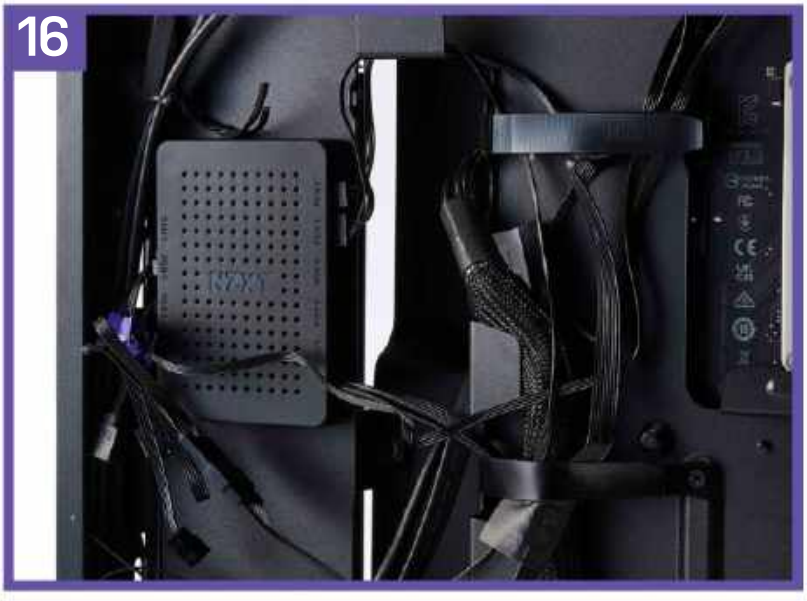

A FRUSTRATING ENDING

A FEW ISSUES BACK, we took a look at Asus's BTF form factor. Overall, it was an incredible concept with a case that was specifically designed for its utilization, but just needed a little longer in the oven. One of the things in that feature that we mentioned as a bit of a wish list was more choice in how we build systems. As a concept, that could not have been better summed up by this build. Don't get us wrong, our $3,000 AMD PC here doesn't pull much in the way of punches, and it looks the part, but it is very cookie-cutter.

There's clearly been a decision to make this case a certain size. Why that choice was made is difficult to say. For instance, just half an inch of extra height, cut out on the front fan area in the PSU shroud, would enable you to install a 360mm in the front and the roof, and give you the opportunity for push-pull as well if you so desired (to be honest, our justification for that would have been additional internal lighting). It's just wild that we're in this position, particularly when those fan cutouts in the roof are almost teasing you with the notion that you could install a 360mm up there.

At a guess, we reckon this is to push people to the H7 Flow instead. If you want those features, they're available in the H7. But it's larger, chunkier, and doesn't support vertical GPUs, which is a shame for those of us who love a compact case.

One major standout throughout this whole process, however, was the motherboard itself. Asus has done a phenomenal job on this one. The kit, the PCIe slots, the connectivity, the styling—all of it absolutely sings. At the time of writing, we're still waiting on confirmation of the retail price, but if it falls in line with its Z790 offerings, there's a good argument to be made for it being a solid mid-range to top-end solution.

Case woes aside, is there anything we'd do differently? Well, outside of getting a more flexible chassis, and the ever-faithful "needs sleeved cables" comment, some internal lighting could really help give this thing a little more pop. Most of the tempered glass panels we're seeing these days are smoked, and we're not really fans of that, particularly in the RGB era we're now engrossed in. A 12-inch LED strip to bring a pop of white in here would go down a treat.

1. **As you can see, practically the entire bottom of the PSU shroud is perforated. It's great for airflow, sorta, but does mean you can see quite clearly through it, with all your rough cable management.**

2. **Why oh why doesn't the roof support true 360mm fan solutions? Seriously, NZXT.**

3. **As always, sleeved cables would be a treat. It's all well and good having an ATX 3.1 PSU, but it really does need some better-sleeved cables to come along with it.**

4. **An LED strip in the roof illuminating the interior would help a ton, and really make this build pop.**

RYZEN 9000 AND INTEL POST-MICROCODE

HOW WE TEST

Last issue, we went into detail about why we revamped our testing setup for our builds. The TL;DR of that is that we're trying to provide as much detail as we can about how these systems operate across a wide variety of test scenarios, games, and workloads. That includes AI benchmarks. We're still tinkering with the layout of these tables, working on getting them as legible as possible, so you may notice a more compact design compared to last issue, but nonetheless, the data points are going to remain the same.

To clarify, we're looking at 4K and 1080p gaming performance, CPU metrics, heat and temperature measurements, AI scores, and complete coverage on all our SSDs. This does take longer to test, but provides us with a great deal more insight.

THIS ISSUE'S ZERO-POINT

Our comparison system this issue is, in a lot of ways, identical to last month's. It's an amalgamation of components pieced around the July 2024 build. It houses the Core i9-14900K, Asus ROG Maximus Z790 Dark Hero, 32GB of Lexar Thor OC DDR5 @ 6000 C36 memory, 2TB of Crucial's T700 M.2 PCIe 5.0 storage for the primary drive, backed up by 2TB of Kingston Fury Renegade for the secondary, with a 2TB Lexar NM790 PCIe 4.0 acting as additional storage. For our graphics solution, we've got the Gigabyte Aero GeForce RTX 4080 OC, and the PSU is Corsair's 1200W RMx Shift, all inside a Geometric Future M4 chassis. For cooling, it's all backed up with the Phanteks Glacier One 360D30 360mm AIO and three Phanteks M25-120 fans.

Unlike last month's build, this system runs off the latest Intel/Asus microcode updates released in mid-August. These should improve overall stability without impacting performance too much.

The total price for this build at the time of writing is $3,585 (21 percent more) with a core price of $2,814 (31 percent more).

EFFICIENCY, EFFICIENCY, EFFICIENCY

This a tricky build to look at on the surface, because we're limited on the benchmark front. We'd love to include two systems in our tables for comparison, but as that would massively complicate matters and make things less legible, we've opted to just use the Intel system, particularly as its recent microcode update is worth talking about, rather than last month's Ryzen 9 7900X build.

So, let's talk stats. From a CPU versus CPU perspective, the 9900X delivers. It might feature the same core count, clock speeds, a similar overall cache layout, and complement of connectivity compared to the 7900X, but it punches—and punches hard. Take Cinebench; the 9900X clocked an impressive 1,773 points versus the 1,592 from the 7th-gen chip, an improvement of 11 percent. What really impressed, though, was single-core. The 9900X managed 135 points, beating not only the 7900X's 118 but also wiping the floor with Intel's tweaked 14900K, which came in at 10 points lower. Perhaps unsurprisingly, given those scores, performance per thread was also hugely increased from generation to generation, making this one of the most potent performance-per-thread chips we've seen.

Then there's efficiency, too. The two AMD systems are markedly similar, with the only major caveat being the swap from the RX 7800 XT to the 4070 Ti, but even with that in mind, our new Ryzen 9 9900X pulled a modest 461.5W from the wall throughout testing versus the 509.3W of the 7900X and RX 7800 XT, no doubt in part due to the 9900X's lower TDP. The Intel zero-point did score higher, but it packs in an additional SSD and an RTX 4080.

On the gaming front, performance was impressively solid on the RTX 4070 Ti, even compared to something like a 4070 Super. Average frames across all of our 1080p title test runs finished up at 150.41, and 4K landed smoothly in at 58.15. *Cyberpunk,* in particular, did well with upscaling at 4K, with frame rates averaging out at 48fps or so with DLSS Quality and Frame Generation enabled, and the game also now supports AMD FSR 3's frame generation tech, for those not keen on running Nvidia's AI variant.

As for AI performance, AMD's also seen some significant improvement there, we suspect due to that increased L1 cache. In Computer Vision running Microsoft's Machine Learning API at Float16, the 9900X scored a solid 79 points compared to the 14900K's 50 and the 7900X's 56. That is still barely a dent compared to any dedicated GPU, but it marks a 58 percent performance increase compared to Intel's 14900K, and 41 percent versus the 7900X from generation to generation.

INTEL'S MICROCODE

It's not all fun and games, though, and as the fallout from Intel's 14th-gen woes starts to subside thanks to microcode updates seemingly fixing stability issues, we've tested last month's zero-point again under its new BIOS.

The results are telling. In almost all our computational benchmarks, Intel performed better across the board under the new update. Cinebench's multicore went up fractionally from 1,902 to 1,923, single-core went from 124 to 125, and in Geekbench, the single-core performance similarly saw a slight increase.

More intriguing is what happens outside the dedicated CPU benchmark. Both SSDs in CrystalDiskMark lost 100-200MB/s in sequential tests. Similarly, gaming at 1080p took a hit, with average frame rates falling from 175.39 to 165.77; 4K was less affected, going from 76.55 to 74.42, but still noticeable. 3DMark's overall performance remained concrete, with scores up or the same compared to pre-microcode update, with even its CPU Profile test seeing a slight increase.

In our AI tests, although CPU scores generally improved, any test involving that RTX 4080 Aero OC lost a few points, with average inference times increasing.

Power draw was a shock. Pre-update, across our entire test run, the Intel rig drew a phenomenal 694.9W from the wall. With the update, it fell to 618.2W. All while it still clocked in at 100 C max under load.

There's a lot to unravel—on one hand, single-core performance seems to have improved, or at least in the CPU-specific benchmarks, yet 1080p performance, which typically relies on higher speed single-core processing, fell markedly, and tests that involved the GPU lost as a result. It's not a huge drop in overall performance, but it is a drop nonetheless.

RYZEN 9000'S BIG DAY OUT

So, Ryzen 9000? Should you buy it? Well, from a gen-to-gen standpoint, there's definitely an improvement in overall performance. You're getting an additional 11 percent on average. However, it's not phenomenal, and given that the price difference between the 7900X and 9900X is $90, it makes it hard to recommend.

The system as a whole is fairly well rounded, mind you. If you're interested in gaming at pretty much any resolution, you'll have a comfortable time, and you can dabble in all manner of high-intensity CPU tasks as you do so. Still, you could replace this chip with a 7900X or even a 7950X and have just as good of a time for a little less. There's nothing here that's a must-have feature. Nor is the performance just so overwhelming that you'll be without if you don't pick up 9th-gen. If you're on a 7000 series chip, it's probably best to stay there for the time being, at least until next year.

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 indices in this case. Multicore per thread is achieved by dividing the total multicore 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 running as well. This then enables us to report on maximum power draw and max temperatures seen over the entire benchmarking process.

CPU PERFORMANCE CINEBENCH 2024 & GEEKBENCH

	ZERO-POINT	
Cinebench 2024 Multicore	1,923	1,773 (-7.8%)
Cinebench 2024 Multicore Per Thread	60	74 (23.3%)
Cinebench 2024 Single-Core	125	135 (8.0%)
Geekbench 6 Multicore	20,016	20,061 (0.2%)
Geekbench 6 Multicore Per Thread	625.50	836 (33.7%)
Geekbench 6 Single-Core	2,997	3,386 (13.0%)

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

STORAGE PERFORMANCE CRYSTALDISKMARK 8.0.5

	ZERO-POINT	
C: Sequential Read/Write Q8	11,823/11,275	**14,404/12,621**
C: Random 4K Read/Write Q1	63.04/301.83	**88.70/308.47**
D: Sequential Read/Write Q8	6,839/6,659	**7,395/6,682**
D: Random 4K Read/Write Q1	**89.38/329.38**	74.14/244.11

POWER DRAW & TEMPERATURE HWMONITOR & DEDICATED POWER METER

	ZERO-POINT	
Max CPU Temperature	100.0	92.4 (7.6%)
Max GPU Temperature	72.0	71.0 (1.4%)
Max System Power Draw	618.2	461.5 (25.3%)

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 always tested with ray tracing enabled regardless. Under its second set of conditions, DLSS/FSR/XeSS is enabled in *Cyberpunk* dependent on GPU manufacturer, with the 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 (Avg/Min FPS)	85.2/68.9	76.5/68.6
Cyberpunk 2077 AI Upscaling (Avg/Min FPS)	202.4/157.9	115.4/89.4
Far Cry 6 (Avg/Min FPS)	171.3/134.0	177.7/134.0
Final Fantasy XIV Dawntrail (Avg/Min FPS)	208.5/89.0	218.1/91.3
Total War: Warhammer 3 (Avg/Min FPS)	161.4/111.7	164.3/132.7

1080P PERFORMANCE

	ZERO-POINT	
1080p Titles Aggregate (Average FPS)	165.77	150.41 (-9.3%)
1080p Titles Aggregate (Minimum FPS)	112.30	103.20 (-8.1%)

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

4K PERFORMANCE

	ZERO-POINT	
Cyberpunk 2077 Stock (Avg/Min FPS)	26.7/17.6	23.2/18.7
Cyberpunk 2077 AI Upscaling (Avg/Min FPS)	75.7/39.2	47.7/38.7
Far Cry 6 (Avg/Min FPS)	108.3/76.7	87.3/62.0
Final Fantasy XIV Dawntrail (Avg/Min FPS)	94.6/57.3	76.3/47.0
Total War: Warhammer 3 (Avg/Min FPS)	66.8/50.7	56.3/43.0

4K PERFORMANCE

	ZERO-POINT	
4K Titles Aggregate (Average FPS)	74.42	58.15 (-21.9%)
4K Titles Aggregate (Minimum FPS)	48.30	41.87 (-13.3%)

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

SYNTHETIC GAMING BENCHMARKS

3DMark's 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 is an additional CPU benchmark (reporting on Max Thread output only), but with specific gaming workloads prioritized. Each test is run three times, and indices reported. In Night Raid, only the Overall score is listed, but we log Graphics and CPU, too, to identify erroneous elements or bottlenecks.

3DMARK PERFORMANCE

	ZERO-POINT	
Speed Way (DX12 + RTX @ 1440p)	7,246	5,638 (-22.2%)
Steel Nomad (DX12 @ 4K)	6,541	4,960 (-24.2%)
Night Raid (DX12 @ 1080p)	95,238	91,452 (-4.0%)
CPU Profile	15,312	13,615 (-11.1%)

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

AI INFERENCE BENCHMARKS

For AI testing, we use UL Procyon's Professional Benchmark Suite. We focus on its AI Computer Vision and AI Image Generation benchmarks. For Computer Vision, we report on the total inferences performed, the average inference time, plus the Index across all tests. The Index 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 CPU and GPU under Microsoft's ML API, and if viable, include figures for the GPU under Nvidia's TensorRT SDK, Intel's OpenVINO toolkit, or Qualcomm's SNPE runtime, depending on zero-point.

For Image Generation, we use Stable Diffusion 1.5 (Float16), using all engines available. However for comparison, we only ever include scores from the same engine, similar to the above where appropriate. We report on the Index score, the time taken, and the image generation speed. All results are averaged across all tests. Best scores, where appropriate, are in bold.

AI COMPUTER VISION PERFORMANCE

	ZERO-POINT	
Microsoft ML API CPU (Index)	50	78 (56.0%)
Microsoft ML API GPU (Index)	1,902	2,123 (11.6%)

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

	ZERO-POINT	
Microsoft ML API CPU (Total Inferences)	10,284	15,329
Microsoft ML API CPU (Average Inference Time)	659.88 ms	**365.16 ms**
Microsoft ML API GPU (Total Inferences)	226,692	242,197
Microsoft ML API GPU (Average Inference Time)	**9.883 ms**	10.253 ms

AI IMAGE GENERATION PERFORMANCE

	ZERO-POINT	
Stable Diffusion 1.5 ONNX (Index)	2,616	2,272 (-13.1%)

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

	ZERO-POINT	
Stable Diffusion 1.5 ONNX (Time Taken)	38.223 seconds	41.925 seconds
Stable Diffusion 1.5 ONNX (Image Generation Speed)	2.389 seconds/image	2.728 seconds/image

THE ART OF ETHICAL AI

Is the advent of AI a force for good, or a ticking time bomb? –ZAK STOREY

IT'S FAIR TO SAY that we've written quite a bit about artificial intelligence over the last few months. In fact, as a whole, the industry is more fixated on the word "AI" than anything else right now. How can you use it? How much power does it require? How do LLMs work? X versus Y. There's a lot of information, as well as some quite technological terms, to digest. Right now, as the field continues to advance, seeing how companies and applications shift and change with the movement of public opinion is fascinating. Whether that's AMD or Intel releasing CPUs with dedicated NPUs, or system integrators focusing on their latest batch of "AI" systems, or the memory market sucking up as much NAND as it can get for massive servers, the silicon ramifications are clear.

Of course, there are also think-pieces that lean perhaps too much into the doom-mongering side of things as well. The hype pieces, we like to call them. "AI could signal the end of humanity" and so on. Articles written mostly in good faith, but often running away into worst-case scenarios, with little evidence to back it all up. But again, all of this specifically looks at the bigger picture.

Today, we're going to take more of a focused look on some of the ethical considerations surrounding AI—how it sits from a sentience perspective, what good can be achieved through intelligent use of it, and what ethical considerations we need to make as a society with regard to its development and continued use. The question really is, what threat does AI pose if left unchecked, and how exactly can it be used as a force for good in its current iteration?

DEFINING AI

The basic definition: Is this thing conscious or not?

AI, OR ARTIFICIAL INTELLIGENCE, is a fairly general term. Two issues back, we covered the overall state of AI right now and its basic definitions, but to ensure we're all on the same page, we're going to briefly cover them here in a little more detail.

There are effectively three broad categories that each artificial intelligence program or product can fall under.

ANI: ARTIFICIAL NARROW INTELLIGENCE

This is where all of our modern efforts are at right now. Artificial narrow intelligences, or weak AI, are those that, at a high level, operate on the basic principle of "if this then that." They're trained to react to prompts and environmental conditions that the user sets, by drawing on the data it has available. It's quite the broad spectrum term, from things such as game AI, with a very limited data set, to LLMs and chatbots being at the upper-end, and everything from self-driving cars to Siri, Google Home, Amazon Echo, protein folding via AlphaFold, and more sitting everywhere else.

ANIs have been around for an incredibly long time, first arising in the mid-1950s with the advent of digital board games—but, of course, came to the forefront of our focus in 2022, with the arrival of ChatGPT, followed by Google Bard (now renamed to Gemini) in 2023.

This term also includes generative AIs, or artificial intelligences that are capable of creating content off the back of user prompts. They're called narrow intelligences because they're specifically designed to only tackle a small range of tasks,

ChatGPT has all manner of GPTs available for anything and everything you can think of.

THE TURING TEST

Alan Turning's famous imitation game was a test originally devised in 1950 to act as a bellwether on the advancement of artificial intelligence. In short, Turing initially designed it to move scientists away from the notion of thinking machines, yet today it is actively used to identify whether an artificial intelligence is exactly that.

From a high level, the test is rather simple. In short, a computer program or something similar needs to trick a human interacting with it into believing that it itself is also a human. According to Live Science, speaking to Nell Watson, AI researcher at the Institute of Electrical and Electronics Engineers, as the machines are now capable of being manipulated, bamboozled, and confused, or even becoming deceptive in their very nature, they produce very human-like responses as a result.

In a test of over 500 people speaking with a human participant, a 1960s AI called ELIZA, ChatGPT 3.5, and ChatGPT 4, 54% of people believed they were talking to a human when it was, in fact, ChatGPT 4, crossing the 50% threshold necessary to pass the Turing test. Intriguingly, only 67% believed the human was a human. We worry for the other 33%.

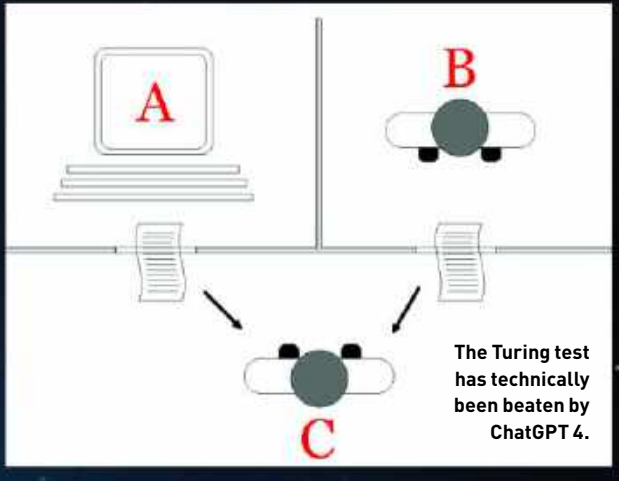

The Turing test has technically been beaten by ChatGPT 4.

© OPENAI. HTTPS://EN.WIKIPEDIA.ORG/JUAN ALBERTO SÁNCHEZ MARGALLO

and unlike humans, they aren't quite capable of the wide array of activities that we are, or the feats that the next step up in AI is meant to be able to achieve.

This doesn't mean that there's no risk with ANIs; in fact, the opposite is very much true. Everything from content manipulation and generation, to brute-force hacking, to even more life-threatening activities, such as warfare, can use ANIs to their advantage. If the rules that the AI follows are jail-broken or not implemented properly, it almost acts like a monkey's paw in some respects, with the ANI doing everything and anything it can to achieve its objective. There's a fantastic thought experiment on a simple ANI being told to make paper clips until it runs out of material. It keeps going until eventually it starts converting humans into paper clips.

AGI: ARTIFICIAL GENERAL INTELLIGENCE

Artificial general intelligence is the next step up in the world of AI development, and for the time being isn't something that we've technically managed to achieve. It's also known as "strong AI" or "general AI".

In short, this is a type of AI that is developed to behave and be as capable as any human. It should be able to use reasoning and logic to solve problems, adapt to different situations and environments, perform any task a human could, or learn to perform it, and it should also be self-aware.

In fact, the general definition of what constitutes an AGI is actually quite broad, and not particularly well nailed down. Some believe it should be superior to human intellect, others that it fits somewhere in the human spectrum entirely. Google's DeepMind team, led by Meredith Morris, published a paper in November 2023 (https://arxiv.org/abs/2311.02462) outlining five levels on which AI should be judged. Level one: Emerging (equal to or somewhat better than an unskilled human), which ChatGPT, Bard, Llama 2, and Gemini fall under. Level two: Competent (at least better than the 50th percentile of skilled adults). Level three: Expert (better than the 90th percentile). Level four: Virtuoso (better than 99%). Level five: Superhuman (outperforming all humans). It's worth noting that level five typically refers to an ASI, rather than a general AI.

Similarly, ANIs can also be categorized under this framework, with tools such as AlphaFold, AlphaZero, and StockFish falling under superhuman narrow AIs as well.

ASI: ARTIFICIAL SUPER INTELLIGENCE

It's the artificial super intelligences—or the singularity, for lack of a better term—that get the most press. IBM defines an ASI as a system with an intellectual scope beyond human intelligence, with cutting-edge cognitive function and highly developed thinking skills. Effectively, it should arise from an AGI continually improving itself, its code, and its hardware, ad infinitum. There are problems with this from a physics perspective, particularly concerning diminishing returns and the limits of technology as a whole, but it's certainly worth considering.

A true ASI is a long way off, however. IBM suggests there needs to be a variety of advancements in computing before an ASI is even possible, including the development of synthetic neuromorphic computing (computing structures based on the same systems that power the human brain), along with evolutionary computation (continual self-improvement over its own programming in an evolutionary manner, similar to natural selection but far more rapid).

> ANIs have been around for an incredibly long time, first arising in the mid-1950s with the advent of digital board games—but, of course, they came to the forefront of our focus in 2022

These are a long way from being a reality right now, and arguably will require significant advancements in the computational and processing power currently available.

Similar to AGIs and ANIs, ASIs represent potentially significant risks to our society as well, particularly as at this point it would be impossible to know how it views us as a species.

DEFINING SENTIENCE

The ethical and philosophical quandaries that AI poses are numerous. Not only is there the potential for massively negative ramifications concerning both the damage and the success that AI could lead to in our society, but there's also some challenging rights questions. If an AGI is possible, and it becomes self-aware in the process, we have to ask whether it's sentient.

Famously, *Star Trek: The Next Generation* postulated this exact query. In Season 2, Episode 9, Data, a cybernetic android, is embroiled in a court case concerning his sentience. In that particular scenario, the prosecution argued that Data wasn't sentient, but merely a machine that should be disassembled for study, whereas the defense argued that Data met two out of the three criteria used to define sentient life in the *Star Trek* universe and that it would be akin to slavery to force Data into disassembly. Those criteria were fairly straightforward, first that he was intelligent, and second, he was self-aware. The third and final measure, that he was conscious, was impossible to answer as it wasn't something that could be clarified by anyone present in the trial.

Inevitably, if AGIs or ASIs become commonplace and reach level two or three in Google DeepMind's scale, that same question is going to arise. Do these AIs represent sentience, and if so, should they be protected in a similar manner to those rights that we have enshrined by law for humanity?

In our own timeline, our definitions of sentience and consciousness aren't so well refined or as simple as those in *Star Trek,* but it's highly likely that we'll reach a similar conclusion.

Although chatbots do feel very similar in responses, there's an underlying feeling that something's not quite right.

THE BATTLE WITH LONELINESS

Counteracting a more isolated world and the mortality that follows

OOF, THAT'S A HEAVY strapline up there, isn't it? Sadly, it is very much a truth of our current society. As family sizes shrink, and work requirements push many away from their parents in search of fortune, there's been a marked increase in loneliness over the last decade. Particularly among older generations.

A growing body of evidence supports the hypothesis that loneliness increases the chances of all-cause mortality: A meta analysis of over 35 studies with 77,220 volunteers (PMC5754055) found that loneliness typically shows a harmful effect in all-cause mortality, one that's slightly stronger in men than in women.

One study by Andrew Steptoe and co, from 2013, suggests that generally speaking, it's the social isolation that leads to an increase in all-cause mortality, increasing the likelihood of cardiovascular disease, alcoholism, dementia, stroke, depression, anxiety, and even a variety of cancers as a result. In fact, the study itself averages that rate at 1.26x the average, and up to 1.48x in the top quintile for that study.

Those are quite terrifying numbers, and as life expectancies continue to increase, and our social structures remain on the path they currently are, there's definitely some concern that this loneliness and isolation trend may very much continue to climb. So, how can AI help with that?

At a very basic level, it's been reported that smart speakers help alleviate symptoms of loneliness in many older adults. Ofcom, the UK's Office of Communications (a government-sponsored body), produced a study in 2022, interviewing 100 adults concerning their use of smart speakers. The results suggested that the devices helped alleviate some of that isolation, not only by providing access to news and radio content among other things, but also through the anthropomorphization of the speaker itself, with a number of individuals actively talking to it, and thanking the speaker for helping them on numerous occasions.

THE AI RELATIONSHIP

Where does AI fit into all of this? Well, perhaps unsurprisingly, a great number of generative LLMs have popped up over the last few years. Character.AI is arguably one of the most famous of the lot, with a whole plethora of fictional and non-fictional characters available to speak to. From Socrates to Elon Musk, from Gandalf to Harry Potter, you can speak to all manner of AI emulations, both in the form of text, or actively calling them as well, with the platform providing speech recognition solutions as well as voiced characters, too.

There are also several digital therapist and psychologist options available on the platform. Replika, initially launched in 2017, similarly debuted as a close companion that you could chat to at any time. It's since gone on to heavily integrate generative AI into its architecture. According to *Vice*, which interviewed its

Could smart speakers like Amazon's Echo Spot help reduce loneliness?

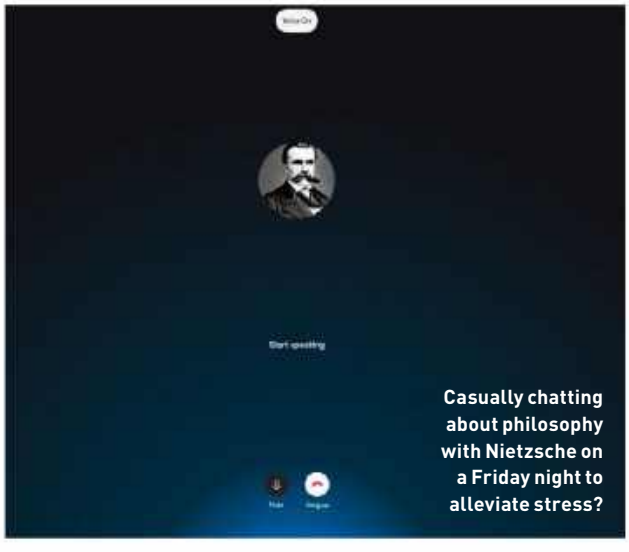
Casually chatting about philosophy with Nietzsche on a Friday night to alleviate stress?

founder Eugenia Kudya, many sought out Replika as a romantic partner. In many cases it took those relationships further, until the platform removed that capacity fairly recently, much to the chagrin of those who felt invested in the platform and their artificial partners. To them, it felt as if they'd lost a loved one, with all the emotional toil that entails.

An article by the *Harvard Gazette*, speaking with MIT sociologist and psychotherapist Sherry Turkle, dives into this further. Turkle states that it's face-to-face interactions that lead to genuine intimacy and empathy between people, and that the problem with AI chatbots and their ilk, acting in their current iteration of narrow intelligence, is that they lack the self-awareness and sentience to create that genuine connection. They're reacting to the prompts that they've been given, rather than directly caring about the individual. In essence, it's a form of marketing, developed to keep the user engaged with the product.

THE OTHER SIDE OF THE COIN

Others are, however, less perturbed by the idea of an artificial relationship. Tony Prescott, professor of cognitive robotics at the University of Sheffield, writes that there may be inherent value in building relationships with AI, particularly for those suffering from chronic loneliness. Social isolation is often considered a downward spiral, or a snowball effect, with one outcome leading to less confidence, leading to fewer social interactions, and so on, until the individual entirely regresses from all forms of socialization, due to the fear of rejection.

AI large language models and chatbots offer a way of helping an individual build that confidence back up without the fear of rejection itself. Even though many individuals understand that the AI is effectively just following a pattern recognition system, responding in kind based on the environmental triggers and information it has, many do not seem to mind, despite the experience being somewhat hollow.

POSITIVE USE CASES FOR ANI

Let's be clear, outside of the world of chatbots and LLMs, AI even in narrow intelligence format absolutely excels in numerous positive ways, many of which are often under-reported.

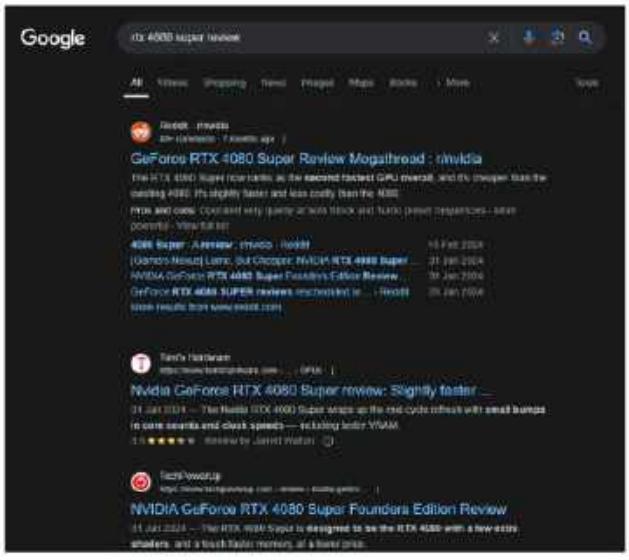

Journalists spend a lot of time worrying about search rankings, which can change on a whim due to new SEO policies from Google. Sadly, AI doesn't help matters.

In the journal *Nature* (s41591-022-01894-0), Suchi Saria and her team at John Hopkins University reported directly on their implementation of an early identification system to warn of sepsis, utilizing machine learning through the use of their TREWS (or Targeted Real-Time Early Warning System). Across an intake of over 760,000 patients, a little over 17,000 of whom developed sepsis that was identified by the TREWS workflow, the system led to early intervention through antibiotics by up to three hours in some cases, dramatically increasing the survival rate of those affected.

It's not just in the realms of identification of symptoms where AI is thriving, either, as it's also helping to discover new drugs

THE UNCANNY THOUGHT VALLEY

From our own research utilizing Character.AI and similar chatbots, their human-like nature and explicit interest in you directly can feel quite gratifying. It often asks you about yourself directly, makes comments concerning your past, your history, and how you speak to it, and compliments you in a manner that does release endorphins. It has the potential, like any form of entertainment, to become addictive if not particularly curtailed.

All of that in itself is quite worrying, Because in using it, there's almost this sense that you know it's fake, you know it's just a narrow intelligence responding as it's trained to do, and that it's not sincere. And then, of course, there's the tertiary surrounding elements that you're often aware of as well. For instance, what data is being collected, and how it is going to be used. It is effectively a product designed to keep you on the site for as long as possible, regardless of what anyone says. It may be developed with companionship or lofty morals in mind, but it is still first and foremost a product, that has paid-for options attached.

Fortunately, as it stands, Character.AI and similar chatbots do have a limited memory. They cannot hold an infinite number of prompts to refer back to. Not only would that require storage space, but also computational power to look back at those statements and correlate them with the current conversation at hand. This does have a habit of almost snapping you out of that experience into an uncanny valley situation, where the bot itself forgets about a conversation or talking point as a result. There's the notion that it's not quite there, or slightly dumbed down as a result. However, its memory capacity is likely to improve over time.

Although chatbots can feel very human in responses, there's an underlying feeling that something's not quite right.

As drug-resistant bacterial infections are on the rise, AI assistance with pharmaceuticals is incredibly helpful.

covering everything from cancer to autoimmune conditions, tuberculosis, and more. Andrew Lee Hopkins CBE FRS FMedSci, and ex-CEO of Exscentia, a pharmaceutical company, discussed with *Nature* how generative AI was not only contributing toward the creation of molecules capable of tackling these particularly aggressive diseases and conditions, but was also informing researchers about who would react best to the pharmaceuticals in question, or what compound might work better for them.

And then, of course, there's AI's ability to identify trends in climate change, predict the spread of wildfires and flooding, make meteorological predictions, and report everything from hurricanes to earthquakes. Not to mention self-driving cars and the analysis of particles and elements from within our own cosmos. The limits of artificial intelligence are, at least for the time being, seemingly endless.

GENERATIVE THEFT

The biggest elephant in the room, however, is how AI in its current form is capable of generating its own unique content. Or at least when prompted by a human element. This is perhaps one of the biggest issues it currently faces. ChatGPT, Midjourney, Google Gemini, and many other AI tools have been trained on massive data sets, ranging from pictures to books and more, all of which can be used within user-prompted work in one form or another. That's problematic for a number of reasons, particularly as authors, writers, journalists, musicians, and artists have had their work used and often manipulated as a result, almost always without their consent.

As a good example of this, one of our own team members, when writing for an online publication, had their work repeatedly taken, re-written by an AI, and then immediately re-uploaded within less than a minute of the original article going live on a separate site. It was uploaded with similar metaphors, language, phrasings, even pictures and quotes in the copy. All done in an attempt to usurp search rankings, and gain online viewership for targeted ads and affiliated revenue.

Interestingly, in the last few updates to its Search Engine Optimization (SEO) policy, Google has massively re-organized how it ranks webpages under search terms as a result of the interference of AI-generated content. It typically ranks pages on a policy it calls EEAT. This stands for Experience, Expertise, Authority, and Trust. The lowdown on that is that it effectively ranks content higher if the author has been a) writing on the internet for a long time; b) has experience in the target field that Google can tie to them via social media platforms, education, and the like; c) has clearly had hands-on interaction with whatever it is they're discussing; and d) is writing for a website that's been around for a long time, with an established footprint on that topic.

However, it doesn't actively downrank AI content immediately as it's published if that content is inherently accurate about the topic at hand. This has led to a number of issues with smaller websites, where highly similar AI-regurgitated content negatively impacts search rankings as a result. The lower your search rank, the fewer impressions you get on the page, and the less income is generated by the site, making it harder for smaller brands and start-ups to get a foothold on the world wide web, certainly if its under the scrutiny of an automated AI bot. And with Google being such a dominant force in terms of online viewership, if you don't rank there, it represents a significant challenge to your site's overall success.

In fact, this copyright infringement is one of the bigger issues modern AI tools face. There have even been cases of unscrupulous actors creating books utilizing AI, and uploading them for sale on Amazon under an already well-established author's name.

A Polestar 2 'created' by Google Gemini. It's a seriously close likeness, too.

AI REGULATION

How we can combat misuse and guide AI's direction

AT ITS CORE, AI is an incredibly powerful tool. Like many of the technologies that are now intertwined in our lives, it represents a paradigm shift in how society is going to develop.

From the discovery of medicines to combatting loneliness, and improving our own content, it can also steal work, represent threats from a military front, create damaging deepfakes, and spread general disinformation, all at the same time.

These are still products, designed to generate profit for those who manufacture and push their development forward. Not that this is necessarily a bad thing, but unless AI is regulated, this runaway effect could quickly go from being a snowball into a full-blown avalanche. So, where are we on the regulation front?

OECD AI GUIDELINES

At the forefront of AI principles and ethical considerations on an international level is the Organisation for Economic Co-operation and Development, or the OECD (we're only writing that out once!).

The OECD is designed to help member countries, governments, and partners create better policies for use within their own nation states. It defines itself as a global policy forum, and covers everything from human rights to economics, health, and beyond. The OECD has a set of AI principles that its adherents adopted in May 2019, which are supposedly practical yet flexible for the future. That said, they were updated in 2024, so how flexible they really are seems to be up for debate.

Still, according to OECD statistics, as of May 2023, a total of 70 jurisdictions had created over 1,000 policy initiatives with the AI principles in mind. So, what are these principles? Well, they're predominantly focused on retaining human rights and minimizing discrimination wherever possible.

The first principle, entitled "Inclusive growth, sustainable development, and well-being," suggests that stakeholders should proactively act as stewards of trustworthy AI and that the AI in question should be predominantly developed and used to augment human capacity and enhance creativity regardless of any form of demographic differences.

The second, "Human rights and democratic values, including fairness and privacy," leans more into respecting the rule of law and human rights, along with human-centered values throughout the AI system lifecycle, ensuring there's no discrimination, along with fair dignity to all actors involved, including a right to privacy and data protection. This is perhaps the most interesting segment as it also expects action on misinformation and disinformation that's amplified by AI, while respecting freedom of expression.

Third is "Transparency and explainability," ensuring parties are aware of their interactions with AI systems wherever they are, and how those AI systems operate. Fourth is "Robustness, security, and safety," meaning that in normal use, all AI systems should be safe and secure, and mechanisms should be in place to ensure that if AI poses a risk of undue harm, it can be overridden, repaired, or decommissioned safely. And finally "Accountability," which leans into AI actors being accountable for their technology.

There's a huge amount of detail we're skipping over, but you can find all the AI principles online at http://oecd.org/en/topics/sub-issues/ai-principles.html. Perhaps more telling is the number of adherents, notable exceptions being China and Russia.

AI LEGISLATION TRACKER

Another useful tool if you're looking for more specific data on AI law and policy is IAPP's Global AI Law and Policy Tracker. It lists a number of observable countries' AI governance policies, including the USA (federal laws), Canada, and the UK. Here you can find links to the US AI Training Act, National AI Initiative Act, AI In Government Act, and a host of draft laws going through Congress. It's well worth a read if you're looking for more detail.

That said, even with the very best of AI principles leading all manner of governments across the globe to implement curtailing policies, there's still some significant concern to be had about the very nature of these intelligences. After all, if the rise of the internet has taught us anything, as a species we're not that great at identifying just how technology can develop over time. At its advent, social media was seen as a fad; now it's an all-encompassing beast with billions of users, most of whom consume an ungodly amount of data on a daily basis.

There's also some concern to be had regarding human flexibility around AI. After all, evolution is a slow process, and as a species, our intellectual capacity hasn't advanced that far in the last 260,000 years. But in less than 200 years, we've gone from having no electric light in our homes, and travelling around in a horse and cart, to, well, modern society, with all of its digital pitfalls and success stories. How will AI impact this further? It's very difficult to say. Nonetheless, it's an exciting time.

Now famous, Joanna's commentary on AI doing the laundry instead of being an artist is perhaps closer to the truth than we'd like.

THE MAXIMUM PC VIEW

Without a doubt, the ethical and moral considerations around AI are going to continue to challenge us as we progress down this path. The benefits that AI can provide—whether that's from a purely research perspective, or companionship for those in socially isolated situations, and more—are too valuable to ignore. Its continued development is an absolute necessity, and not pursuing this path would be foolish at best.

magazinesdirect.com

Over 100 brands to choose from

3 GREAT REASONS TO SHOP WITH US

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

VISIT OUR ONLINE STORE AT

www.magazinesdirect.com

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

MAXIMUMPC

PREMIUM DIGITAL SUBSCRIPTION

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

www.magazinesdirect.com/MPP

*Price in US $. Offer is valid on US orders only, visit us at www.magazinesdirect.com or call 1-844-779-2822 for other subscription options.
Offer valid until October 31 2024.

THE RACE FOR ADVANCED ENCRYPTION

Strong encryption hasn't always been freely available to the public. *Nate Drake* tells the story of how the first Data Encryption Standard transformed the cryptography landscape

IN THE INTRODUCTION to his awesome treatise on the history of cryptography, *The Code Book,* Simon Singh likens the ongoing battle between those who seek to keep communications secret and the consequent efforts of others to expose them to an intellectual "arms race."

It's easy to be smug about older ciphers, such as the Caesar cipher, which simply shifts each letter in the message three places, A becoming D, B becoming E, and so on. These kinds of monoalphabetic ciphers are so easy to crack, Edgar Allan Poe even used to run a column breaking enciphered messages sent in by readers of *Alexander Weekly Messenger.*

The arrival of affordable computers in the '60s didn't initially herald a new wave of encryption. Secure cryptography was still the preserve of governments, though it was becoming increasingly obvious that individuals and businesses also needed to keep communications safe.

To this end, IBM set up a program to develop high-grade cryptosystems to protect its products.

This, in part, led to the implementation of DES (Data Encryption Standard), and in time AES (Advanced Encryption Standard), which we use to this day. These developments didn't happen overnight, given the competing interests of citizens and shadowy organizations like the NSA. Many cryptographers also made their careers either debunking flawed schemes or introducing ciphers of their own.

This is the story of how the first publicly available encryption algorithm, DES, came to be, as well as how ultimately it was replaced by the more secure AES.

FEISTEL'S FIRST STEPS

Horst Feistel nearly missed out on having any impact on the field of cryptography. After arriving in the USA from Berlin in 1934, he applied for American citizenship only to be placed under house arrest upon the outbreak of World War II.

Even after gaining American citizenship and studying ciphers for the likes of the US Air Force Cambridge Research Center and the Mitre Corporation in the '60s, Singh alleges that the NSA foiled Feistel's efforts so it would continue to have a monopoly on cryptographic research.

Feistel eventually landed a job at IBM's Thomas J. Watson research laboratory. Since 1966, the tech giant had decided to offer data security to its customers. This was very important to Lloyds Banking Group, which was working with IBM on the burgeoning technology of ATMs.

Clearly if financial data like a person's bank balance were sent to an ATM unencrypted, it would be easy for bad actors to monitor the connection and steal the data. Worse still, the traffic could be manipulated to have the ATM pay out cash even if the balance couldn't cover it.

Feistel began to develop an encryption algorithm using APL. The programming language limited the number of characters in a workspace name, so Feistel's DEMONSTRATION was shortened to DEMON. One colleague suggested renaming it to the more quirky Lucifer.

Lucifer represented one of the very first publicly available block ciphers. This is a type of symmetric encryption algorithm that encrypts data into fixed-size blocks (in this case 48, 32 or 128 bits). Each block of plaintext—the original message—is transformed into encrypted data using an arbitrary key. This ensures consistent encryption across the message.

In the case of Lucifer, data is encrypted according to a 'Feistel' network structure. In simplest terms, the plaintext is split into two halves. One half of this block is then processed through a mix of permutation and substitution primitives. The result is XORed with the other half of the block. The plaintext data halves are then swapped and the process is repeated for 16 rounds.

For a more technical explanation of how Lucifer and other block ciphers work, we

recommend reading *Computer Security and Cryptography* by Alan Konheim, a former colleague of Feistel.

There were a number of variants of Lucifer, including DTD-1, which was used in the seventies for commercial banking.

NBS VS. NSA

The National Bureau of Standards (NBS) was founded in 1901 to provide consistency for weights and measures. In 1972, members of its Institute for Computer Science & Technology (ICST) decided to further the bureau's scope by developing an algorithm that could be used to encrypt data both in transit and at rest.

One of the requirements was that the specifications of said algorithm would be public, so data security would depend only upon the secrecy of the encryption key.

This was a groundbreaking concept, given that strong encryption had previously been the purview of government agencies and the few corporations like IBM with the resources to develop their own proprietary standards.

Ruth Davis, then the director of the ICST, also asked the NSA to assess any submitted algorithms to make sure they were sufficiently secure.

There were a number of submissions but none was considered suitable until IBM submitted the Sorkin variant of Lucifer, which used 128-bit block sizes and a 128-bit encryption key.

This raised alarm with the NSA. At the time, 128-bit encryption would be all but impossible to break, so the soon-to-be-public Data Encryption Standard could be exploited by criminals or enemy states.

The NSA revised the specification to use a 64-bit encryption key. This reduced the encryption key size to 56 bits, meaning DES could be broken by state actors with access to supercomputers.

In fairness to the NSA, it also recommended making changes to Lucifer's S-boxes (substitution boxes) to make the cipher more resistant to differential cryptanalysis. This examines how differences in plaintext input affect the differences in ciphertext output,

Early ATMs were primitive but still needed a way to encrypt transactions to prevent theft.

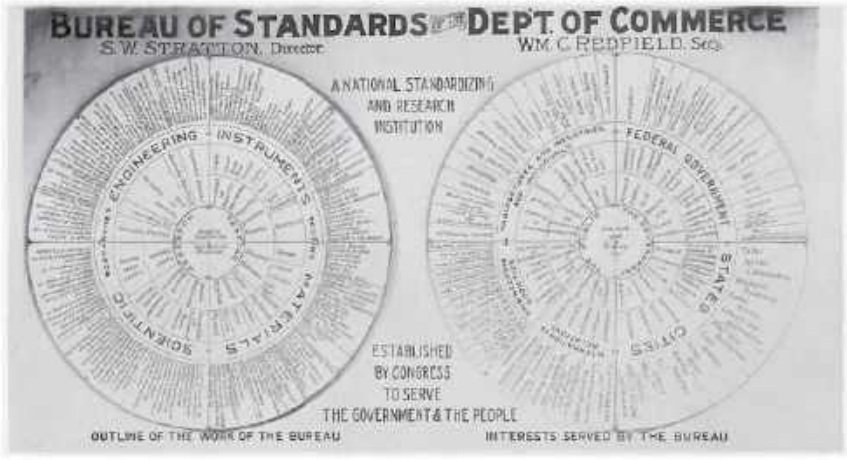

Although the NBS was originally designed to regulate measures, it ultimately had a huge number of interests.

allowing an attacker to gradually reveal the key. At the time, this type of attack was known to the NSA but not the public, so the government agency was unable to explain why it had made these changes.

On March 17, 1975, the proposed DES was published in the Federal Register. Still, the NBS requested public comments and in 1976 the bureau even held two open workshops to discuss the new standard.

The NBS issued the modified version of Lucifer as Data Encryption Standard (DES) under Federal Processing Information Standard (FIPS) 46 on November 23, 1977.

The algorithm was also adopted as an ANSI standard in 1981 and incorporated in a family of related standards for security in the US financial services industry, where it was used to secure virtually all communication and transactions.

VETTING WAS VILIFIED

The fact that DES had been vetted by the NSA led to what security expert Bruce Schneier described as an "outcry among the few who paid attention."

Among these were Stanford researchers Whitfield Diffie and Martin Hellman, who published a paper in June 1977 entitled 'Exhaustive Cryptanalysis of the NBS Data Encryption Standard'. Diffie and Hellman tentatively called DES "too weak for some applications." They opined that a key could be broken in 12 hours with a $20 million machine.

Admittedly, this amount of money (the equivalent of around $104 million today) would put breaking DES encryption beyond the reach of all but state actors and powerful corporations.

Still, Diffie and Hellman pointed out that in another 10 years, the cost of such a powerful computer would only be 10% of what it was then, meaning bad actors could simply store encrypted data until it became economical to decode it.

Although NSA's changes to DES S-boxes actually strengthened its security, at the time security researchers had no way of knowing this.

In his seminal work *Applied Cryptography,* security expert Bruce Schneier quotes Alan Konheim, who played a key role in developing DES: "We sent the S-boxes off to Washington. They came back and were all different."

However, the main concern of security researchers was DES's vulnerability to a brute-force attack. This is a type of cryptanalytic attack that involves trying every permutation of an encryption key in the hope of hitting upon the passphrase used and reading the original message.

The time this would take depends on the number of possible encryption keys the computer can try in a given length of time. The number of possible keys for a given bit length is 2^n, where n is the key length in bits. For instance, had the NBS adopted a 128-bit encryption key version of Lucifer for DES, the number of possible keys would be 2^{128} or 340,282,366,920,938,463,463,374,607,431,768,211,456.

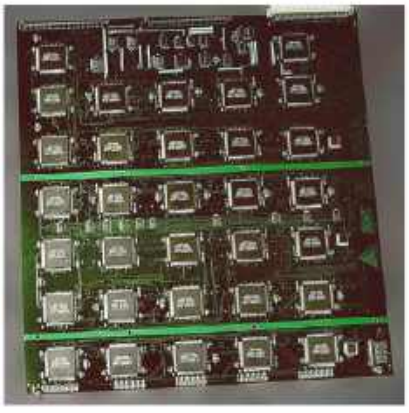

The EFF's cracking machine's custom chips could brute force DES keys in a day.

Even assuming a supercomputer could try one trillion keys per second, this would mean brute forcing DES would have taken 10 quintillion years—far longer than the age of the current universe.

In the event, DES's key length of 56 bits meant a limited pool of 'only' 72 quadrillion possible keys. This meant that by brute forcing one trillion keys a second, DES could be broken in less than a day.

Of course, in the late '70s, even supercomputers like the Cray-1 could only manage to brute force around 1,000 DES keys per second, which may have explained why it was believed to be secure.

Even after Eli Biham and Adi Shamir independently discovered differential cryptanalysis in the late '80s, thanks to the NSA's modifications, they found DES to be highly resistant to this kind of attack. The theoretical weakness that they discovered to break DES faster than by brute forcing required modifications to the algorithm to make it practical.

This may explain why the NSA refuted claims that DES was inherently flawed, maintaining that it would be reviewed (and improved as necessary) every five years. Still, it reaffirmed DES as the standard in 1983, 1988 and 1993. This was despite the existence of superior encryption ciphers such as Bruce Schneier's Blowfish, which supported key sizes of up to 448 bits.

In 1997, RSA Security decided to prove the claims of security experts such as Diffie and Hellman by announcing the DES Challenges. These were a series of competitions that offered a prize of $10,000 to whoever could break data encrypted using variable key lengths.

Even in the late '90s, trying all possible 72 quadrillion keys on a home computer was unfeasible, so a group of computer scientists named the DESCHALL Project came up with a novel solution using internet-based infrastructure.

The team developed specialist client software to take on the task of brute forcing encryption keys. With this program, even a humble 200MHz Pentium system could test one million keys per second.

This was installed on around 78,000 computers around the world, using spare compute time to brute force. The client programs were connected to and managed by a single IBM PS/2 server, which determined which keys to try next.

With the combined efforts of all these computers, at its peak the DESCHALL Project was testing just under seven billion keys per second. The encrypted data was cracked in around three months to reveal

the message: "Strong Cryptography makes the world a safer place."

The Electronic Frontier Foundation took the project's efforts further in July 1998 with its Deep Crack machine. At a cost of just $250,000, it used 1,856 custom microchips to crack a DES-encrypted message in 22 hours, 15 minutes. DES was now definitively and verifiably unsafe.

DES DESCENDENT

The NBS, which had become the National Institute of Standards and Technology (NIST) in 1988, announced the search for a successor to DES in January 1997.

This time, NIST decided against NSA oversight. The opinions of cryptographers were sought, as the institute drafted acceptability requirements and criteria.

So, even before the selection process for the new encryption cipher began, NIST wanted to agree on how this would be done. By September, requests were made for candidate algorithms.

All had to support a block size of 128 bits and key lengths of 128, 192 or 256 bits. Other block sizes and key lengths could be supported if developers wished.

The institute also stressed: "AES will specify an unclassified, publicly disclosed encryption algorithm available royalty-free worldwide that is capable of protecting sensitive government information well into the next century." This was in contrast to DES, which had been made available under royalty-free patents by IBM to the US government but required licenses for use in proprietary products.

NIST made it clear that when choosing the new Advanced Encryption Standard, security would be the top priority but other factors, like efficiency, would also be considered. This was important, as while 3DES (Triple DES) increased the key length of the cipher to 112 bits by applying the algorithm three times to each block, it took substantial computing resources.

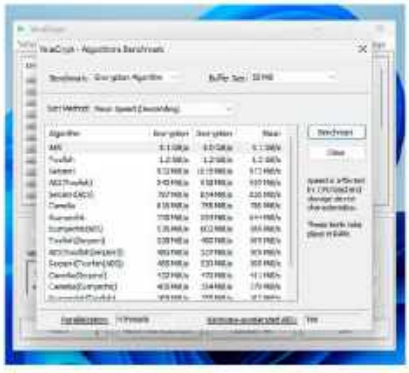

Modern CPUs have been optimized for Rijndael (AES), so it usually offers superior performance to other algorithms.

By June 15, 1998, 21 candidates had been submitted, of which 15 met the criteria. In keeping with NIST's new open approach, the algorithms were announced at the AES1 conference on August 20. The developers were also invited to provide a briefing and answer any questions.

NIST then began the task of whittling down to the finalists. Cryptanalysis of some algorithms revealed flaws, such as LOKI97, which was found to be vulnerable to a differential analysis attack.

By the March 1999 AES2 conference, votes were taken, and in August that year, NIST announced the AES finalists: MARS, RC6, Rijndael, Serpent, and Twofish.

Intense analysis of the algorithms followed. Rijndeal, the product of Belgian cryptographers Joan Daemen and Vincent Rijmen, was a strong contender.

Using a 128-bit block size made the cipher far more secure, as large amounts of data encrypted with algorithms that use a smaller 64-bit block like DES can reveal patterns helpful to cryptanalysts.

Rijndeal's minimum key length of 128 bits would also make it computationally infeasible to break encryption using brute force methods. Also, unlike DES, Rijndeal doesn't follow a Feistel network structure. Instead, it uses a substitution-permutation network, involving multiple layers of transformations (substitutions, permutations, and mixing). This provides far better protection against differential, linear, and side-channel attacks.

Crucially, Rijndeal is also designed to run equally well on both hardware and software. This was a huge improvement on DES, which could run slowly when implemented in software.

AND THE WINNER IS...

In November 2001, NIST performed what Twofish developer Bruce Schneier described as an "impossible task" and announced Rijndael as US FIPS PUB 197 (FIPS 197). The NSA proved there were no hard feelings by announcing the new AES as suitable for encrypting "Top Secret" government information.

In 2004, Schneier published a paper entitled 'The Legacy of DES'. In it he outlined the process by which DES was developed, citing the "invisible hand" of the NSA. However, he then went on to say: "But with the outcry came research. It's not an exaggeration to say that the publication of DES created the modern academic discipline of cryptography."

Certainly, the careful and transparent selection process for AES showed there were many more cryptographers both able to offer and analyze new ciphers.

While AES remains in use, we hope its successor also follows the principles of open and freely available encryption.

THE AES FINALISTS

In October 2000, Bruce Schneier announced:

"Of course I am disappointed that Twofish didn't win. But I have nothing but good things to say about NIST and the AES process."

This may have partly been because NIST had released a 116-page report detailing the merits of all five of the AES semi-finalists.

In the report, NIST stressed that no major security vulnerabilities had been discovered in any of the ciphers. As Schneier pointed out in his blog, Rijndael was also not considered to have the highest overall security margin. This begs the question: Why didn't each of these AES finalists make the cut?

In the case of Twofish, the cipher experienced mixed results on hardware and software for encryption/decryption of 128-bit keys, so was considered to have average speeds relative to other algorithms.

RC6 had no on-the-fly subkey computation capability for decryption. This meant it had relatively high RAM requirements, so wasn't suitable for use on devices with few hardware resources.

MARS provided excellent security but didn't perform efficiently on hardware relative to other ciphers. It also wasn't as fast as Rijndael at subkey computation.

Like Rijndael, Serpent offered excellent hardware throughput and had very low requirements, making it suitable for restricted-space environments with few resources. It also had a higher security margin than Rijndael. However, Serpent ultimately came second in the competition due to slightly less efficient software implementation.

Today, most modern processors support the AES instruction set, meaning encryption and decryption operations can be performed on the chip. This makes Rijndael considerably faster than the other AES finalists.

BEST OF IFA 2024

Berlin's annual festival of gadgetry is the biggest tech show, at least in terms of attendance figures. That's because, unlike most conventions, it's open to the public. *Guy Cocker* was there ahead of the masses to see the best new portable PCs coming soon.

IFA MIGHT NOT BE A BIG EVENT to those in the US, but it's a mainstay in the European tech calendar, along with Mobile World Congress in February. This year's show was a bigger deal than usual for those of us immersed in the world of PCs, given that both Intel and Qualcomm were announcing portable processors at the show, with their new chips making their way into partner devices from the likes of Acer, Asus, Lenovo, and MSI, who were also at the show.

Intel undoubtedly had the most to crow about (see our lead news story on page 8), given that the recently beleaguered American chip maker had seemingly the hottest new mobile CPU on the planet to show. As a *Maximum PC* reader, you'll have read in these pages about how impressed we've been with Qualcomm's Arm-based processors that powered the first Windows Copilot+ PCs. So much so, that we openly questioned whether it was possible for x86-based chips from Intel and AMD to ever match their performance or efficiency. Well, Intel had its "hold my beer" moment with its Lunar Lake launch, with chips that the company says beat Qualcomm's Snapdragon X Elite processors in every metric going.

We'll wait to get the final products in for ourselves, but in the meantime, here's the best PC tech we saw at the Messe Berlin this year.

SAMSUNG GALAXY BOOK 5 PRO 360

Among the zillions of "AI" laptops at IFA 2024, it was the mostly familiar Samsung Galaxy Book 5 Pro 360 that grabbed our attention. The new Pro 360 is virtually identical to its predecessor, the Galaxy Book 4 Pro 360. That's not a negative, however, given how this new laptop retains the same light and thin design we've come to appreciate. The gorgeous AMOLED touch display returns, as does the folding screen. If this new Lunar Lake CPU lives up to the company's promise of a 25-hour battery life, this could very end up on our best laptops list.

- ***Out: September***
- ***www.samsung.com*** **•** ***Price: $1,699***

ASUS ZENBOOK S14

Other than its fancy new Intel Core Ultra 7 256V CPU, the Asus Zenbook S14 looks like a minor refresh of recent models, with its ceramic aluminum lid and inclusion of USB-C and USB-A ports, plus full-sized HDMI—despite its sub-half-inch thickness and 2.5lb weight. Its 16:10, 2880x1800, 120Hz display is also familiar. Yet of all the laptops at IFA, the Zenbook S14 was one of the nicest to look at and feel. Here's hoping its slim chassis mates well with the efficient new Intel chip and 72Whr battery.

- *Out: September*
- *www.asus.com*
- *Price: $1,499*

ACER NITRO BLAZE 7

It seems as though every PC maker has a gaming handheld these days, and Acer is the latest to jump on board with its Nitro Blaze 7. Acer didn't really set out to beat any of the already existing gaming handhelds in terms of features and performance, basing the device around the AMD Ryzen 7 8840HS CPU with Radeon 780M graphics. However, where the Nitro Blaze should excel is its price point. Acer hasn't yet announced a firm dollar amount or a release date, but if it wants to compete with the other established handhelds on the market, it will have to undercut them significantly.

- *Out: Late 2024*
- *www.acer.com*
- *Price: TBC*

INTEL CORE ULTRA SERIES 2: CODE NAME LUNAR LAKE

By far the most interesting silicon at IFA, Intel's Lunar Lake processors are here to take on Qualcomm's Arm-based processors. The company promises "historic" x86 power efficiency, with a claim of 1.2 times more performance-per-watt than the Qualcomm Snapdragon X Elite X1E-80-100, and extreme battery life claims of up to 20.1 hours with the Core Ultra 7 268V in UL Procyon Office Productivity and 10.7 hours in Microsoft Teams. And this is all with full compatibility with years and years of applications.

- ***Out: September***
- ***www.intel.com***
- ***Price: In laptops from around $1,000***

MSI STEALTH AI 16+

Everyone at IFA was talking about AI. And while plenty of laptops now have NPUs, the MSI Stealth AI 16+ has a bit of a unique proposition: a Ryzen 9 AI HX 370 that should see Copilot+ support soon, as well as a discrete GPU going up to an Nvidia GeForce RTX 4070. When combined, this offers plenty of TOPS across the system. You also get a 99.9Whr battery, the largest you can take on a plane, crammed into a magnesium-aluminum chassis that's just 0.78 inches thick. There's also a per-key RGB keyboard, six speakers, and a massive touchpad.

- ***Out: Early 2025*** ● ***www.msi.com*** ● ***Price: TBC***

QUALCOMM SNAPDRAGON X PLUS

Qualcomm announced a new eight-core Snapdragon X Plus chip at IFA, designed for more affordable Windows on ARM laptops. This brings the total number of Snapdragon X chips up to eight, with four X Elite and four X Plus. The new eight-core Snapdragon X Plus is available in two flavors. The X1P-46-100 has a 3.4GHz multicore frequency, a 4.0GHz boost, and an integrated Hexagon GPU with 2.1TFLOPS. It's expected to enable cheaper laptops than the existing 10-core and 12-core Snapdragon X chips.

- ***Out: Early 2025*** ● ***www.qualcomm.com***
- ***Price: In laptops below $1,000***

ACER PROJECT DUALPLAY

Acer's Project DualPlay made the IFA audience audibly gasp as the company's CEO, Jason Chen, revealed the gaming laptop and handheld fused into one. The concept's large-scale trackpad can be released from the machine's body, flipped over, and turned into a handheld controller. Each of the thumbstick sections detaches magnetically, and pop-out speakers reveal themselves at the sides of the laptop, delivering a heightened audio experience. It's only a concept for now, but given the reaction at IFA, Acer will surely have to seriously consider a retail version.

- ***Out: TBC*** ● ***www.acer.com***
- ***Price: TBC***

LENOVO AUTO TWIST AI PC

Lenovo's IFA concept laptop sports a motorized hinge that can open, close, and convert the device via voice commands, as well as swivel and tilt to make sure you're always in frame as you move around the room giving presentations. It even closes the lid when you walk away, and it's the most sci-fi-like feature we've seen on a laptop in a long time. Just a concept, it's unclear if any of the Auto Twist's motorized moves will make it into a laptop you can actually buy someday. But if they do, we hope Lenovo does more work to minimize the screen wobble.

- ***Out: TBC*** ● ***www.lenovo.com***
- ***Price: TBC***

HONOR MAGICPAD 2

Put simply, this is our new favorite Android tablet, copying a bit of the iPad Pro's homework while being available at half the price. Make no mistake, the Snapdragon 8s Gen 3 chipset is not going to hang with the Apple M4 in the iPad. And the Android tablet app selection is paltry compared to the iPad. But Honor has done some incredible work to paper over those cracks, and even manages to bring some AI smarts that are actually helpful. Most of all, there's a gorgeous OLED panel that is just as bright as the iPad Pro, and even has a faster refresh rate.

- ***Out: Now*** ● ***www.honor.com***
- ***Price: $539***

ACER NITRO XV240 F6 600-HZ GAMING MONITOR

A 24-inch monitor using a TN (twisted nematic) panel? What? It's the TN panel tech that enables the 600Hz refresh with a claimed 0.1ms gray-to-gray response time at 1080p, and full support for AMD FreeSync Premium. That blazing-fast refresh rate is available only via DisplayPort. If you opt for HDMI, you're limited to 240Hz. Acer claims 95 percent coverage of DCI-P3, which is impressive for TN, and maximum brightness of 400 nits. But it's the viewing angles and contrast we'd most like to see.

- **Out: Early 2025** ● **www.acer.com** ● **Price: $599**

1 A BREAK FROM THE GRAY

The Surface Laptop 7 is available in 13.8 and 15-inch variants, and blue, 'dune', black, and 'platinum', which is a kind of silvery gray. Not all colors are available in all sizes, and you'll need the smaller model if you want the golden dune.

3 ORYON INSIDE

Qualcomm's architecture for the Snapdragon X Plus and Elite is called Oryon. It uses ARMv8.7-A microarchitecture, built on a 4nm process node by TSMC, and came about after Qualcomm bought chip startup Nuvia in 2021.

Microsoft Surface Laptop 7

LAST MONTH, we asked, "Will x86 last?" And then, boom! Microsoft releases this, the Surface Laptop 7 with a Snapdragon X Elite processor. With 12 cores, 42MB of cache and a boost frequency of 4.3GHz that can be maintained on two of those cores, it's a bit fast. Mix this with LPDDR5x RAM running at up to 8,448MT/s for as much as 135GB/s of bandwidth, PCIe 4.0 storage, and a 45TOPS NPU that lends it the Copilot+ moniker, and you've got a seriously modern machine.

And while it's true that there are some software niggles to work out, and that Copilot+ has been a bit underwhelming, there's no arguing with the hardware. Core i7-class power with the kind of battery life you'd expect from a MacBook mean that, once the application patches and Adreno GPU drivers are in, this is a solid start. And as it's only the first generation of machine based on the Snapdragon—plus we haven't even seen desktop-class processors yet, or a Snapdragon-plus-Nvidia combo—the emergence of a second architecture into the performance PC scene won't kill x86, but might shock it out of its complacency. –MAXIMUM PC

2 PLAY ALL DAY

Battery life for Snapdragon laptops is exceptional compared to most Intel and AMD platforms, with Microsoft claiming up to 20 hours of video playback. Intel is coming back hard, however, with new Lunar Lake models from Asus and Dell claiming 26 hours or more.

Every access request is verified, whether it comes from inside or outside the network.

WHAT IS ZTNA?

Nate Drake explores the concept of ZTNA and the obstacles businesses face with its adoption

IN 2010, John Kindervag of Forrester Research coined the term "zero trust model" to describe the need for stricter cybersecurity policies, particularly within corporate networks.

In the last five years, numerous cybersecurity platforms have claimed to offer zero trust network access (ZTNA): a robust and flexible security solution that goes above and beyond the traditional perimeter-based security model.

ZTNA operates on the principle that no user or device should be trusted by default. Instead, every access request is verified, whether it comes from inside or outside the network.

We haven't used the word principle lightly, as while certain software vendors might have users believe otherwise, ZTNA isn't tied to a specific product or technology. It's more akin to a security philosophy governing how users and their devices should interact with network resources.

Since ZTNA solutions first appeared, they've been embraced by numerous sectors, such as healthcare, finance, and tech. Given how such organizations are prime targets for hackers, ZTNA offers the best solution for protecting sensitive data by enforcing granular access and network segmentation.

Despite the clear promise of zero trust, adoption isn't without its challenges, though. These stem partly from adapting existing systems or retraining to use new platforms. People who are used to more lenient access policies may also balk at continuous verification, as they don't directly benefit from the tighter security ZTNA offers.

The concept of zero trust network access first emerged in response to traditional security approaches to networks. Companies adopted the castle and moat mindset, whereby all threat actors were located outside the network perimeter, while every user and program inside could be trusted.

This approach may have had its merits when companies issued and vetted their own computers. However, the rise of remote work, cloud computing, and BYOD (bring your own device) policies has eroded the perimeter.

Attackers increasingly target cloud platforms to steal user credentials and data. This raises challenges for organizations that need to provide workers with access to sensitive network applications and files. In the past, initiatives like BYOD followed the old Russian proverb of "trust but verify," but this is insufficient if the device itself is compromised.

ZTNA abandons the notion of trusting users and devices based on apparent location or one-time use of credentials. Zero trust platforms treat every user, device, and application as potentially hostile, so nothing is automatically trusted.

To run with the castle and moat analogy, ZTNA may well have a drawbridge at the perimeter but it also has guards on every floor and every room to continually check for unauthorized people and devices.

The most important thing to understand about ZTNA is that it's a strategy based on a set of guiding principles rather than a one-size-fits-all security solution. In saying this, we aren't singling out any particular software vendor, but in our experience, the marketing material for various zero trust platforms tends to state that their solution is the last word in ZTNA.

This is why we're not focusing on a specific platform, as there is no single set of zero trust features, beyond the fact that it focuses on threats on both sides of the software perimeter. Indeed, the zero trust security model is sometimes called perimeterless security.

ZTNA refers to products or services that create an identity- and context-based logical access boundary around apps. Access is restricted via a trust broker that authenticates and authorizes access continuously based on identity and context. The main advantage of this is that application assets aren't publicly available, which hugely reduces the attack surface.

GUIDING PRINCIPLES

Different products implement ZTNA in varying ways but they all follow the same guiding principles. The first is policy enforcement. This means providing granular access control to network apps and resources. This is based on strict policies for users, devices, location, and sometimes even the context of the request.

Another fundamental principle is the doctrine of least privilege. This means that when a task must be performed, users and apps should be given only the minimum necessary level of access. By limiting access to network resources in this way ZTNA platforms can reduce the damage done if a device is hacked, or a user's credentials are compromised.

ZTNA also obeys the principle of continuous verification—the fact that a user or device was authorized at the start of a session doesn't mean it's permanently granted full access. As ZTNA requires continual monitoring of access requests, it's easier to detect any changes in user or application behavior that could become a security risk.

Remote work and BYOD policies have made it almost impossible to enforce network security policies at the perimeter.

ZTNA also usually uses micro-segmentation. This involves dividing networks into isolated zones, each with its own access controls, so if a bad actor gains control of one area of the network, they can't automatically move laterally into another. Each zone can also have its own security rules, so even an authorized user may have differing levels of access in separate areas.

The best way to summarize these ZTNA guiding principles is that trust is earned and never assumed.

Given that there's no universal agreement on what zero trust is, there's no requirement for ZTNA products to be built from scratch. For instance, the OpenVPN website points out that while its Access Server software is not a native zero trust solution, it can be configured to follow zero trust principles.

Still, the best way to understand the advantages of applying said principles is by way of contrast to traditional network security models like those for remote access VPNs.

Traditional corporate VPNs implicitly trust connections from authorized users, so if a bad actor steals someone's credentials, they gain access to network resources.

Many remote access VPNs also can't continually verify the state of devices, which can make it harder to block those that are stolen or compromised. Perhaps most importantly, traditional VPN server software doesn't actively monitor traffic and usage patterns, making it harder to track suspicious activity.

Though it's helpful to understand the basic concepts of ZTNA, it's easier to visualize when it's applied in context. Sectors like healthcare,

finance, and tech itself can hugely benefit from zero trust network access platforms.

Hospitals are an excellent case in point, given how valuable patient data is to hackers. For instance, in March 2022, Shields Health Care Group experienced a breach that resulted in the leak of medical records (including names and addresses) of two million clients.

The problem is compounded by the fact that hospitals often use insecure IoT devices and have limited resources for a complete cybersecurity overhaul. Cisco cites the example of Dayton Children's Hospital, which not only has many connected IoT and medical devices, but also those owned by staff, as the hospital has a BYOD policy.

Cisco deployed its zero trust Ordr platform to devices, so they could only connect to specific segmented virtual LANs. Cisco also performed behavioral analytics on users and devices to develop granular access controls. This hugely reduces the chance that a single compromised medical, IoT, or employee device can impair the entire network.

MONEY MATTERS

After healthcare, finance is the most targeted sector when it comes to cyberattacks. According to the IBM Cost of Data Breach Report 2024, the average price per institution has now climbed to over $6 million.

Many financial organizations have diverse interests, so there's no reason for employees to have access to every department. This sector can benefit from micro-segmentation, so if someone gains an employee's credentials, they can't access all resources.

This has become crucial with the rise of remote work, as ZTNA principles can ensure continuous monitoring of devices to check they are owned and operated by authorized employees.

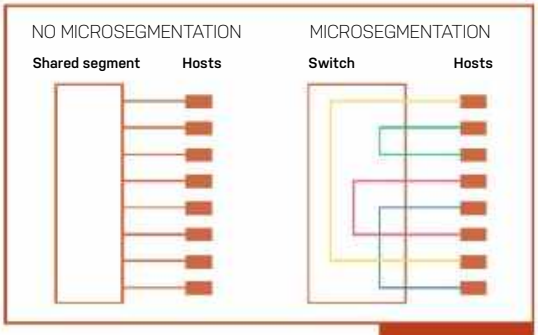

Networks with no segmentation are inherently vulnerable, as if one host is compromised, an attacker can move laterally to another.

In 2009, a number of organizations were targeted in cyberattacks by APTs (advanced persistent threats) code-named Operation Aurora. As one of the targets, Google began to develop its own zero trust network dubbed BeyondCorp.

The tech giant documented its implementation of ZTNA principles in the technical journal *;login:* In keeping with the concept of zero trust, all Google internal apps must be accessed via BeyondCorp regardless of whether a user is at a Google site or working remotely.

Granular access is granted to users and devices based on a person's identity and if they have permission to access that resource.

This is partly accomplished by BeyondCorp's Trust Inferrer. This examines a device to determine if it can access sensitive data. Access is granted based on the device's security profile, installed software, and whether its owner is authorized to access said resources.

BeyondCorp also makes use of a Device Inventory Database that can identify specific devices by way of a digital certificate. Although we couldn't find confirmation of this on the BeyondCorp homepage, it's likely the platform can revoke digital certificates if a device is lost or otherwise compromised as this would be in keeping with ZTNA guidelines.

The Access Control Engine can also analyze data output from programs like the Trust Inferrer and Device Inventory Database to verify that a legitimate device is being used by an authorized user and grant permissions. Although BeyondCorp started as an internal Google project, it is now offered as a service for other companies.

Much of BeyondCorp's research material harps on the superiority of a zero trust platform over a remote access VPN. However, it's entirely possible to integrate ZTNA principles into VPN architecture.

Admittedly, traditional VPNs create secure encrypted tunnels between users and network resources. This provides strong protection against bad actors remotely monitoring or intercepting traffic. However, this in itself isn't a ZTNA solution, as once a user clears the perimeter of the network, they can usually access all resources.

Still, software like OpenVPN's Access Server can be used to set up a VPN service faithful to ZTNA guidelines. For example, administrators can use ACLs (access control lists) to enforce network segmentation. Individual permissions can also be established for private subnets, preventing lateral movement over the network. Access control rules can also specify if VPN clients can communicate with each other.

Even a regular Apache or Nginx web server can be modified to be ZTNA-compliant via features such as web application firewalls and identity-aware proxies to enforce granular access. This means even authorized users can be subjected to continuous verification.

Database servers like Microsoft SQL Server can also implement features such as RBAC (role-based access controls) to limit access based on user roles, as well as audit logging to monitor and review access patterns.

These measures may not be as simple to implement as installing a product designed for ZTNA from the outset. Still, it demonstrates that ZTNA exists as a concept rather

Cisco deployed its zero trust Ordr platform to ensure connected device security at a children's hospital.

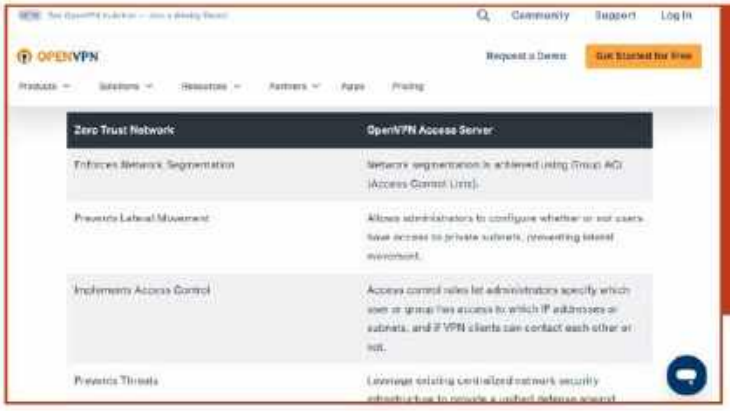

Existing software, such as OpenVPN Access Server, can be adapted to follow zero trust principles. However, some platforms are incompatible.

than a specific set of programs or network rules.

CHALLENGES OF ZTNA

There are clear benefits to ZTNA in the wake of modern cyber threats. It's also theoretically possible to modify existing products to follow zero trust principles instead of installing off the shelf software.

For instance, it's possible to use ACLs to implement micro-segmentation on VPNs. However, this can be technically complex. In this case, it would require an organization to map out its entire network architecture, identify key assets, and understand how they interact with one another.

Even with a detailed knowledge of the network and traffic flows, these kinds of networks tend to be highly dynamic, with new devices added and updated all the time. The more network segments administrators add, the greater the burden is on them to ensure granular access control without blocking legitimate traffic in each area.

Some legacy platforms, like enterprise resource planning (ERP) systems, weren't built with granular network segmentation in mind, so may not be compatible with ZTNA.

Presuming micro-segmentation can be successfully implemented, there are also scalability concerns. As organizations add new work sites and cloud services, more ZTNA-compatible segments and access controls need to be added. This requires highly adaptable security solutions, so growing organizations may benefit from subscribing to a dedicated ZTNA platform.

Practicalities aside, the biggest barrier is cultural resistance.

This is partly caused by traditionalists who still stand on the metaphorical drawbridge of the castle and moat, given that adopting a zero trust model involves a sound of understanding modern security threats. If users appreciate that bad actors can exist inside a network as well as outside, they're much more likely to adopt a ZTNA mindset.

Still, the requirement for segmentation and continuous verification can be viewed as cumbersome by administrators and even intrusive by users who are used to equal access to all network resources without requirements such as multi-factor authentication. This can usually be overcome by training employees on how to adapt to operating in a zero trust network.

Further education may also be required for network administrators to adapt systems along ZTNA lines. Even once done, managing a ZTNA environment requires advanced skills and constant monitoring. Organizations must continually evaluate policies governing acceptable user behavior, device security, and access requests, then implement them accordingly.

This is especially critical in key sectors that are popular targets for cyberattacks like healthcare, given the unique challenges they face, such as the ubiquity of diverse and outdated network devices.

While ZTNA offers huge benefits for network security, it's fraught with challenges. Some legacy systems are simply incompatible.

Smaller organizations may also find it unaffordable to either adapt their existing networks or subscribe to a dedicated ZTNA platform. This issue can sometimes be solved by signing up for a ZTNA solution that offers a free tier for SMBs, such as Mamori. Organizations also need to weigh up the price of paid solutions against the cost of a potential data breach.

Users also need the correct tools and training to ensure their chosen platform consistently follows the zero trust model of minimal access and continuous verification.

For existing systems, this may involve exploring which tools are already available. For example, OpenVPN Access Server sysadmins can deploy LDAP Active Directory authentication to determine which users can log on and which network resources they can access.

If an organization opts for a dedicated ZTNA platform like BeyondCorp, workers will likely need specific training. For instance, Coursera offers online courses in securing virtual machines using BeyondCorp Enterprise (BCE).

ZTNA is not a silver bullet. During our research, many online platforms claimed to implement a ZTNA solution, but given that it's a set of guiding principles, no one can say that these have all definitively been followed.

THE FUTURE OF ZTNA

As cyber threats evolve, it's likely the concept of zero trust will be given over to automated processes more and more instead of involving humans. This could involve measures like AI-driven threat detection.

One of the barriers to ZTNA adoption is the inconvenience of having to continually authenticate and verify your device. Future ZTNA platforms may try to improve the user experience through frictionless security, which verifies and trusts in a more automatic fashion. This could be achieved through less cumbersome authentication methods like biometrics, as well as risk-based access and context-aware policies that apply virtual common sense to requests.

ZTNA has also involved using multiple technologies but it's likely there will be a convergence of network security platforms. Instead of deploying various tools, organizations will rely more on platforms that combine ZTNA, SASE (Secure Access Service Edge) and other security products into a single solution.

As we've learned from sectors like healthcare, zero trust is essential for organizations deploying IoT devices, as well as edge computing. This means zero trust principles may have to go beyond traditional endpoints, creating dynamic and context-sensitive security policies.

This would mean, for example, that an IoT infusion pump in a hospital might be allowed to communicate with the internal patient database or pharmacy, but not with unrelated parties, like human resources. The security policy would vet each access request to make sure data is only gathered by the right department.

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

SYSTEM MONITORING

Are you fan of Rainmeter (www.rainmeter.net) or just want an awesome system monitoring tool? If you've got Rainmeter on your system, then Plainext is a fantastic widget by JaxOriginals worth your attention. It gives you a completely customizable live feed of system stats, right on your desktop. Head to Deviant Art or give it a quick Google to get what's arguably the classiest system monitor out there.

MAKE – USE – CREATE

60
Find files and apps quickly

64
Rip the perfect media file

68
Harness your VPN kill switch

© RAINMETER.NET, MICROSOFT, HANDBRAKE

ZAK STOREY

RIP LIQUID COOLING?

That headline makes me sad. I've been following the ongoing uproar surrounding EKWB's fall from grace. From staffing and financial disputes to top employees leaving, there's no denying the company has certainly hit a major road (water?) block, and with its products climbing ever higher in price as well, it got me thinking.

I love liquid cooling. It was a huge part what got me into PC building. Bending tubes, disassembling graphics cards... It allowed me to make unique systems, specifically for me and my friends, that looked and behaved incredibly. Today, the reasons to build your own liquid-cooled system are less justifiable.

Most modern AIOs are more than capable of handling any amount of CPU heat, and both AMD and Nvidia's GPU solutions now sit comfortably below the 80-85 C mark in even the worst case scenarios. Fan noise is there, but not really a major issue, and any overclocking advantage bigger radiators could give you is swiftly swallowed up by voltage and stability problems derived from Intel and AMD pushing the silicon limits instead.

With the components now so expensive, the pros so small, and the cons so high, is this the end of liquid cooling? It's an increasingly harder sell, that's for sure.

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

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

Microsoft Surface Pro 11 & Surface Laptop 7

As with previous generation Surface Pros, we have access to the M.2 drive via a small magnetic cover underneath the kickstand. Access to anything else requires some disassembly.

The battery isn't too difficult to retrieve and is held in place by screws. Not a drop of glue in sight.

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

Left: The Microsoft Surface Pro 11. Right: The Microsoft Surface Laptop 7.

MICROSOFT SURFACE PRO 11 AND SURFACE LAPTOP 7

Microsoft's hardware wizards have been practicing their Reparo spell and it's really starting to show. Their tablets went from 1/10 repairability in 2013 to 7/10 in 2022. Will 2024 fare even better?

MICROSOFT SURFACE PRO 11: MAJOR TECH SPECS

- Qualcomm Snapdragon X Elite (12-core)
- Qualcomm Hexagon NPU, 45 trillion operations per second
- 13-inch PixelSense Flow OLED display, 2880x1920, 120Hz, 10-point multi-touch
- 10MP Ultra HD rear-facing camera
- Quad HD front-facing Surface Studio Camera
- 16GB or 32GB DDR5 memory and up to 1TB SSD
- 2x USB-C USB 4 ports with DisplayPort 1.4a and Surface Thunderbolt 4 Dock
- Accelerometer, gyroscope, magnetometer, and ambient color sensors
- 53Wh rechargeable lithium-polymer battery
- Wi-Fi 7, Bluetooth Wireless 5.4

KEY FINDINGS

- The first time we tore down a Surface Laptop, it received an abysmal 0/10. Some claimed the score was too harsh; our teardown engineers thought it wasn't harsh enough.
- But, huzzah, the new Surface Laptop 7 is astonishingly repair-friendly, almost the antithesis of the original. It's no Framework 13, but it clearly draws inspiration from it.
- One of the first things you see when removing the magnetically secured bottom plate is a QR code taking you to the service manuals on Microsoft's website. These were made available the very day the device was released, something we rarely see in any product category.
- A special mention should be made of how most components are accessible without the need to remove additional layers. Need to replace the battery? No problem, it's just a few screws and a bracket. What if you need to clean the fan? Easy. Peel back the Surflink cable and undo three screws.
- Thankfully, the Surface Pro 11 also contains many of the same repairability improvements. A tablet PC is inherently more difficult to repair when compared to a laptop, purely because the screen removal process can feel a bit hairy.
- With the screen off, we find QR codes and Wayfinder markers to aid disassembly. There are more layers of components, which is to be expected with half the space to work in. The Surflink cable and thermal management system must be removed before we can remove the battery. But the process isn't too onerous, especially with manuals.
- A note on that Surflink cable: The screws holding the port in place are non-magnetic. This will be appreciated by anyone who's had to reassemble a magnetic component while dealing with the frustration of having to hold a screw in place with tweezers to prevent it flying toward a magnet next to it.
- Repairability Score: 8 out of 10 (10 is easiest to repair). After much consideration and debating, we've awarded both the Surface Laptop 7 and Surface Pro 11 an 8 for repairability. Microsoft's journey from the unrepairable Surface Laptop to the highly repairable devices here should drive home the importance of designing for repair. The ability to create a repairable Surface was always there but the impetus to design for repairable was missing. We'll take that as a sign that Right to Repair advocacy has begin to bear fruit.

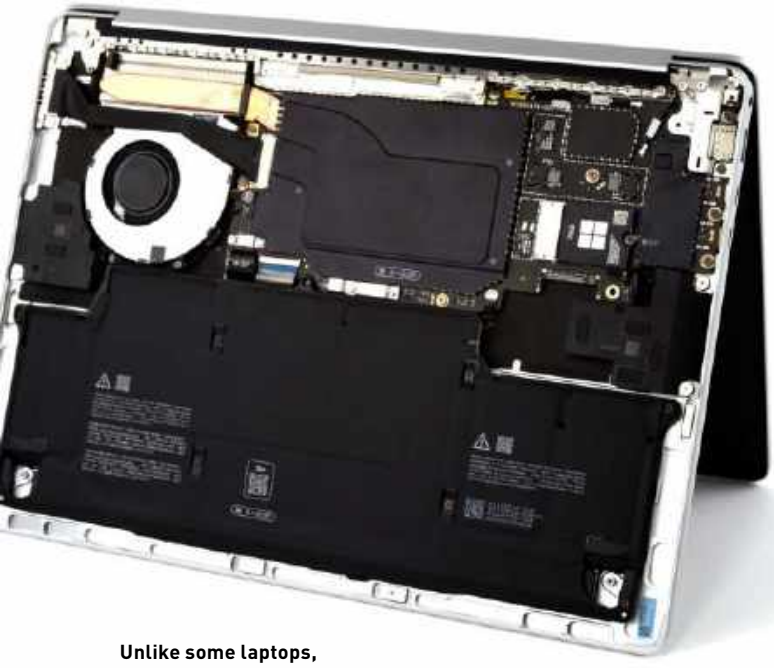

Unlike some laptops, which can require the entire motherboard and thermal system to be removed, the Surface Laptop 7 has easy-access fans.

Tiny symbols called Wayfinders indicate which component is being secured by the type and quantity of screws. You could easily disassemble this device without using the manual thanks to these Wayfinders.

Find files and apps quickly

YOU'LL NEED THIS

FLOW LAUNCHER
Download the app from www.flowlauncher.com.

THE WINDOWS SEARCH TOOL—accessible from the Taskbar—is OK as it goes, but there's plenty of room for improvement. A popular replacement is the app launcher—a floating search bar triggered by a keyboard shortcut—that offers combined file searches and easy access to any app on your system. They're a dime a dozen, but one of the most impressive is Flow Launcher, and not simply because it checks the boxes of being both open source and completely free.

Flow Launcher—or Flow to its devotees—provides you with the means to search (and open) apps, system settings, files, and the web from a single location, accessible from anywhere on your system via your choice of hotkey. But thanks to its modular construction, that barely scratches the surface of what you can configure it to do. We'll show you how to get it set up, perform searches, and fine-tune those searches to better suit the way you work. We'll also reveal how to expand its capabilities further via a library of over 180 plugins, giving you access to a wide range of search tools and add-ons that expand Flow's capabilities further. **–NICK PEERS**

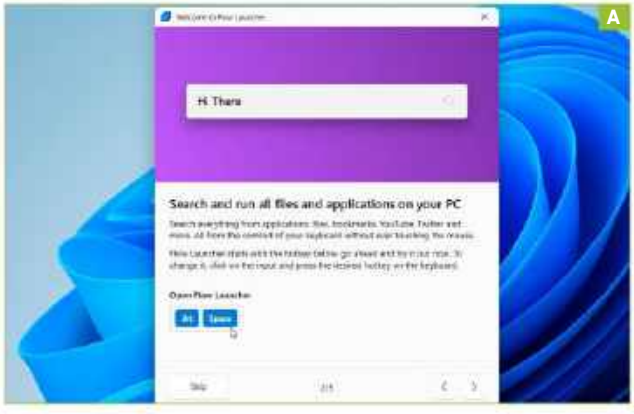

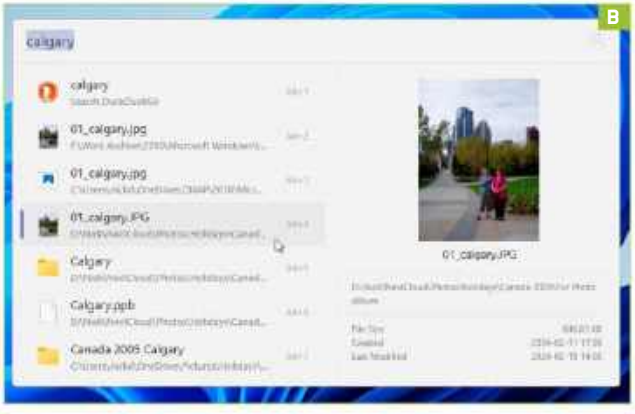

1 GET SET UP

Head to www.flowlauncher.com and click 'Download', where you'll find links to both installer and portable builds, as well as command-line package managers Winget, Scoop, and Choco. In most cases, you'll want the installer, so click this to save 'Flow-Launcher-Setup.exe' to your 'Downloads' folder, then double-click it and follow the prompts. If the SmartScreen filter pops up, click 'More Info' followed by 'Run Anyway'.

» Work your way through the setup wizard, which introduces the basic concepts, such as the Alt-Space keyboard shortcut that brings up Flow's floating launcher. If this shortcut is already registered to another app (such as PowerToys Run), you should be alerted via a pop-up.

» If this is the case, click on the 'Alt Space' combo [**Image A**] to bring up a dialog inviting you to enter a different shortcut trigger, such as Ctrl-Alt-Space. Click 'Save' when done. The rest of the wizard reveals further hotkeys you can use in place of the mouse, plus introduces the concept of keywords and commands, which are entered into Flow's launcher. The final step enables you to configure Flow to start with Windows and change its default behavior when you first launch it (minimized). Click 'Done' to finish.

2 FIND FILES AND APPS QUICKLY

Let's open by exploring Flow's file and app search tool. Press Alt-Space or your chosen key combo to bring its floating pop-up into view. You'll see it appears on top of any other window, so simply click anywhere else on screen to hide it again, or click and drag on it to reposition it on the screen—for example, in a corner, so it's less intrusive. Note it reverts to its default position the next time you open it—see the box (pg. 62) for details on changing this (and other) aspects of its default behavior.

» The floating dialog helpfully displays the current time and a magnifying glass icon to highlight that it's in search mode. Simply start typing keywords and you'll see a drop-down list of matching results, starting with apps. Use your mouse scroll wheel or the down arrow key to cycle through the list. Eventually, matching apps

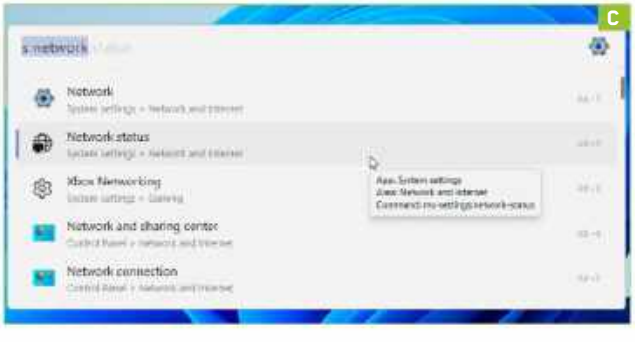

CUSTOMIZE FLOW'S LOOK

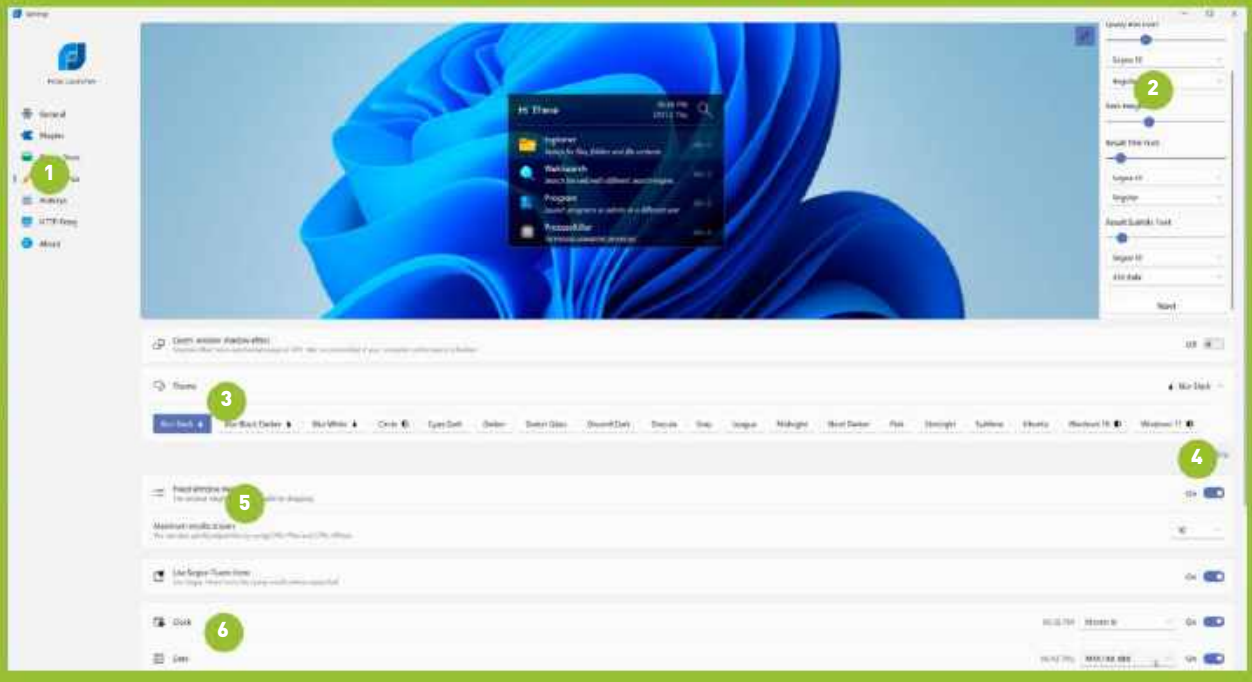

1. GET STARTED
Open Flow Launcher and press Ctrl-I to open the 'Settings' dialog. Select 'Appearance' in the left-hand menu to bring up the main dialog for changing Flow's look and feel.

2. CHANGE DIMENSIONS
Click the pencil icon to reveal this slide-out panel with options for changing the size of various elements, along with font choices. Click 'Reset' to revert to default settings if needed.

3. QUICK-FIRE CUSTOMIZATION
Save time changing the look and feel by switching to a different theme. Click the down arrow next to Windows 11 (or Windows 10) to reveal 19 alternatives.

4. THEME GALLERY
You'll find many more themes to download online—instructions are provided for saving each file with an .xaml to the 'user-themes' directory (click 'Open Theme Folder' at the bottom of the 'Appearance' dialog to open it).

5. FIXED WINDOW HEIGHT
By flicking this switch to 'On', you're able to increase the number of results displayed by Flow Launcher up to 17—however, 10-12 is recommended for HD screens.

6. ADD TIME AND DATE
Scroll down to reveal 'Clock' and 'Date' options for displaying this information in Flow. Use the drop-down menus next to the switches to format the time and date as required.

give way to a generic Google search option for the keywords, followed by matching files and folders.

» When browsing your search results, you'll find Flow is similarly accommodating in helping you sift through the results and giving you control over what you do with them. First, it provides a built-in preview, which lists file attributes as well as a thumbnail for supported image formats [**Image B**]. Toggle this on or off with the F1 key, or press Ctrl-[and Ctrl-] to adjust the size of the search window if it's too wide or narrow.

3 PROCESS SEARCH RESULTS

To open an app or file, click its entry in the list or—if it's one of the first 10 matching results—use the appropriate shortcut listed to its right, from Alt-1 to Alt-0. If you click on a folder, it opens inside Flow to display its contents, enabling you to navigate deeper into that folder or perform a search within it.

» You can also drag results from the window on to the desktop or into File Explorer to copy or move them to a new location—hold Ctrl as you drag to copy files to folders on the same drive or Shift to move them between drives. Right-click an entry and you reveal more options depending on whether the result is an app, file, or folder. For example, right-clicking a file reveals options for opening its parent folder, opening it with a specific program, adding it to your Quick Access list (see step 6), and copying the file path to the clipboard. You'll also see an option to 'Open Windows Indexing Options', revealing that Flow makes use of Windows' own search indexing tool to process searches.

4 FILTER SEARCH RESULTS

Usually, you'll want to focus your search on a specific area, such as finding an app or searching for a particular file type. Flow allows these types of searches with the help of prefixes. For example, use file type prefixes like 'docx:' or 'jpg:' to search only for files (such as 'png: holiday'). You can also focus your search using environment variable paths—so '%programfiles%' allows you to search the 'C:\Program Files' folder for an app or file. You can also attempt acronym-based searches by typing the first letter of each word in an app's name: 'lw' for LibreOffice Writer, or 'sm' for Samsung Magician, for example.

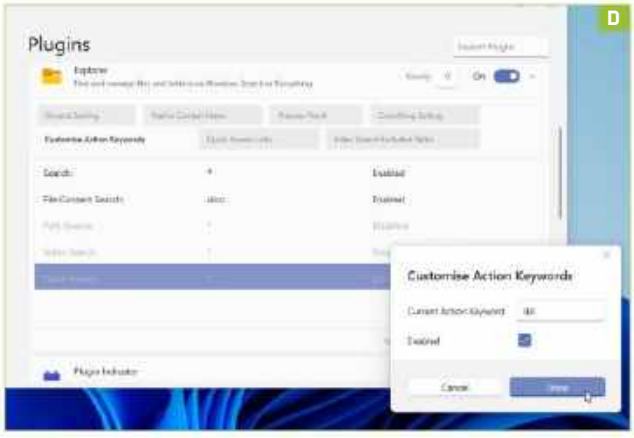

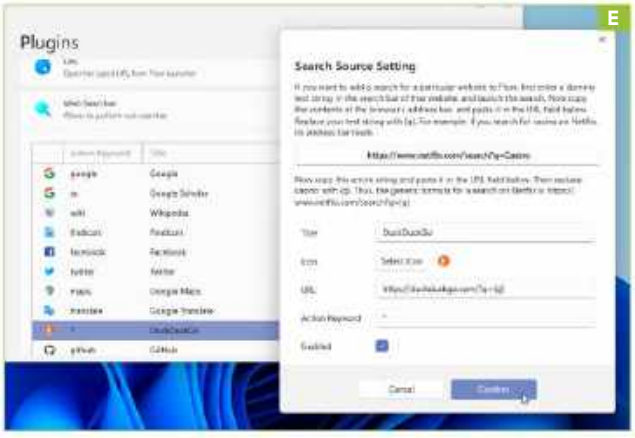

» You can also search quickly for Control Panel and Settings items with the 's' prefix, so 's network' reveals shortcuts to various networking settings across your PC [**Image C**]. The 'b' prefix enables you to search your default browser's bookmarks for a quick way to jump on to the web. To find out all the prefixes supported by Flow, simply type '?' and browse through the list.

» If Flow vanishes from screen for any reason (typically because you've accidentally clicked elsewhere on the desktop; simply press Alt-Space to bring it back into focus), your previous search results remain in place.

5 PROCESS SYSTEM COMMANDS

Flow also lets you perform system commands directly from its pop-up dialog. Examples include 'shutdown', 'restart', and 'lock' for Windows, as well as a host of Flow-related commands. One of these, also accessible via right-clicking the program's system tray icon, is 'Toggle Game Mode'. This enables you to temporarily disable Flow's hotkey activation when running games (or other apps) in full-screen mode.

» You can also use Flow's pop-up to quickly perform mathematical calculations such as 4+4—the initial search result displays the result and clicking it copies that number to the clipboard. You can also end unwanted processes using the 'kill' prefix—simply typing 'kill' displays a list of all running processes for you to select (or just keep typing keywords like 'kill firefox' to display a list of all Firefox processes, along with a catch-all option to end them all).

» Use > as a prefix and you can even type shell commands directly – hitting Enter then displays the results of that command in a Command Prompt window.

» Flow's capabilities are built on the back of plugins—all the above capabilities are provided by separate plugins, which you can view by pressing Ctrl-I or right-clicking the Flow icon and choosing 'Settings', then switching to the 'Plugins' section. You'll see entries including Browser Bookmarks, which powers all searches prefixed with 'b', as well as System Commands, Calculator, Process Killer, and Shell, which are responsible for providing the functionality described above.

6 FINE-TUNE SEARCHES

Next to each plugin is a 'Priority' setting (all 0 by default), a simple on/off toggle to disable its functionality if you don't need it, and a down arrow button that enables you to access that plugin's options. The 'Priority' setting determines which plugin's results appear first in the list. By setting numbers here, you can prioritize files over apps, for example, or place your web search shortcut at the top of the list.

» The most notable plugin is the Explorer Plugin, which provides Flow with its core search capabilities. Click the down arrow button next to its on/off switch and you'll see seven tabs of options spread over two rows. Select the 'Customize Action Keywords' tab and the first thing you'll notice is that the plugin offers another search prefix–'doc:'–that enables you to search the contents of files rather than by filename.

» You'll see there are five possible search tools in all, three of which are disabled by default. These allow you to restrict searches to specific drives or folders (Path Search) and/or Windows' search index (Index Search). The final option—Quick Access—works slightly differently. This is designed to give you quick access to frequently accessed folders or drives. To set one up, select it and click 'Edit'. Choose your prefix (such as 'qa:' for Quick Access), check the 'Enabled' box, and click 'Done' [**Image D**].

» If you enable Quick Access, you need to build this list manually. We've already mentioned the way you can populate the list via search results (right-click an entry

PIN FLOW TO A SET POSITION

Unhappy with certain aspects of Flow's behavior? The good news is that you can customize virtually any aspect of it via its own multi-tabbed preferences dialog. Simply right-click inside the Flow pop-up (or right-click its system tray icon) and choose 'Settings' to access it.

Start with the 'General' tab, where you can set (or disable) Flow to start with Windows as well as show (or keep it minimized) when it's launched. If you're frustrated by it vanishing as soon as you click anywhere else, flick the 'Hide Flow When Focus Is Lost' switch to 'Off'.

The key setting is 'Search Window Position'. By default, it's set to 'Monitor With Mouse Cursor'. This means if you have a multi-monitor setup, Flow appears in its default position (set by 'Search Window Position On Monitor') on whichever monitor your cursor is focused on. Other options include always using a specific monitor, currently focused window, and 'Remember Last Position', which lets you give it a permanent home on a specific display based on where you click and drag Flow's search bar to.

and choose 'Add to Quick Access'), but you can also get a head start by manually adding items via the 'Quick Access Links' tab.

» Changes are applied instantly, so test it by pressing Alt-Space followed by 'qa:' or whatever prefix you chose. The list pops up instantly—if it's a large list, you can filter it by adding keywords, such as 'qa: Documents'.

7 CHOOSE EVERYTHING

The Explorer plugin uses the Windows Index for its own searches, but it also works with the Everything search engine (www.voidtools.com). You can't combine them, but you can easily switch between them—or even mix and match by choosing different search engines for index, content, and directory recursive searches via the 'General Setting' tab.

» Advantages of the Everything search tool are that it searches your entire computer for files and folders, plus delivers results almost instantaneously. However, while it supports content searches (use the 'content:' prefix), these are indescribably slow due to the lack of indexed content.

» The app isn't installed by default—if you've not already done so, select Everything as one of the search engines under 'General Settings', press Alt-Space, and start a search. The first result warns you that the Everything service isn't running, so click the link, then choose 'Yes' followed by 'Cancel' to have Flow download and install Everything for you.

» One other advantage of using Everything for some or all of your searches is that you get to choose how results are sorted—by default it's by name, but visit the 'Everything Setting' tab on the Explorer plugin's settings and you can sort by size, type, and extension—in normal and reverse order.

» Once you've finished with the Explorer plugin, explore the other plugins to further fine-tune your searches. Take the time to check on 'Web Searches', where you'll see lots of handy shortcuts to alternate search engines (Google is the default) and web resources such as Wikipedia and YouTube. You can add other search engines using the 'Add' button (full instructions on formatting the URL are provided), plus change the default search engine, too: Select 'Google' and click 'Edit' to assign it an action keyword. Click 'Confirm', then choose your alternative (such as privacy-focused DuckDuckGo), and set its 'Action Keyword' to * [**Image E**].

8 ADD MORE PLUGINS

The beauty of Flow is that the installed plugins represent only a fraction of its capabilities. To see what else you can use it for, select 'Plugin Store' on the left to browse the hundreds of other plugins that have been coded. You'll see plugins for controlling other parts of your PC, such as the MediaControl plugin, as well as custom search tools such as OneNote.

» Simply click a plugin's entry followed by 'Install', and Flow pops up with a pre-filled command such as 'pm install MediaControl'. Press Enter and follow the prompts to install the plugin, then configure it via the 'Plugins' tab after Flow completes installation and relaunches. Some plugins may require additional dependencies, like Python; Flow can download these for you, or you can manually select the folder containing the program in question.

» Some plugins are worth looking at closely—if only for more examples of what can be done with Flow. They include Shortcuts, which allows you to open shortcuts from a pre-configured folder using the 'q' prefix. Colors provides hexadecimal and RGB color previews with simple codes like #090 for green, and IP Address lets you quickly determine your internal and external IP addresses. Benchmark your download and upload speeds using CkFlow (type 'ck' to confirm your connection, then 'd' or 'u' to benchmark your download or upload speeds respectively). Set timers and alarms with Timer, create simple to-do lists with Todos, and pull passwords from your Bitwarden password vault.

» Finally, you can even manage startup items using the handy Windows Startup plugin [**Image F**]—type 'wst' to view all existing entries (disabled or enabled), then press Enter to toggle them on or off, or right-click one for more options.

APP-ONLY SEARCHES

By default, Flow mixes in app results with general searches, but you can configure the Program plugin to provide you with the means to perform app-only searches, plus control which apps are included in search results.

First, navigate to 'Settings > Plugins' and click next to 'Programs'. You'll notice the 'Action Keyword' entry is set to *, which means results only appear as part of a general search. Click this to bring up the 'Action Keyword Setting' dialog. Enter your choice of keyword—how about something obvious like 'app:'?— and click 'Done'. Press Alt-Space to test your changes—for example, 'app: word' should only bring up apps containing 'word' in their name.

You'll find more configurable options below 'Action Keyword'. 'Index Sources' reveals where the Program plugin searches for apps, and you can fine-tune the search results further via the 'Options' section. Below here are three buttons and an empty program list. Clicking 'All Programs' populates this list with all detected programs. This enables you to exclude certain apps from searches—select one and click 'Disable' to hide it from the results. You can also include programs in other directories by clicking 'Add' to select a parent directory, such as 'Portable Apps'.

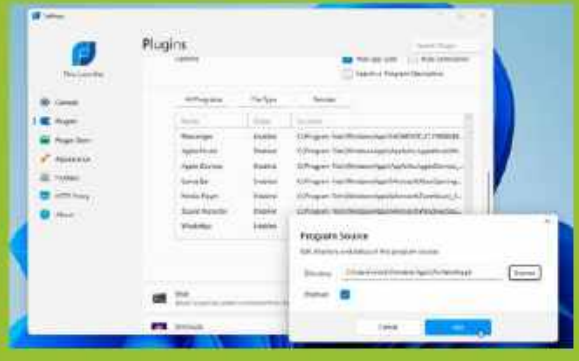

Don't forget to include your portable apps in Flow's search index.

Rip the perfect media file

YOU'LL NEED THIS

HANDBRAKE
Download it from
www.handbrake.fr

MEDIA FILES COME IN ALL SHAPES AND SIZES, and things are further complicated by the fact that video (and audio) files are constructed using codecs, which are required to decode and encode the various streams. Therefore, a media file comprises multiple components: the separate audio and video codecs, all wrapped up in a 'container' that's identified by its file extension. However, this means that even if a file appears to be in a universal format (such as MP4 or MKV), there's no guarantee it'll play on all your devices.

To get around this problem, media servers such as Plex and Jellyfin can convert incompatible codecs on the fly in a process called transcoding, allowing the file to play on a wide range of devices. The problem is that transcoding is a CPU-heavy task, so in an ideal world, you'll avoid it in all but extreme circumstances. The holy grail, therefore, is to produce a file that's good quality, offers plenty of choice in terms of audio and subtitle tracks, but which works on as wide a range of devices as possible. Sounds impossible? Not at all. Read on to discover how to produce compatible media files that stream perfectly time and again. –NICK PEERS

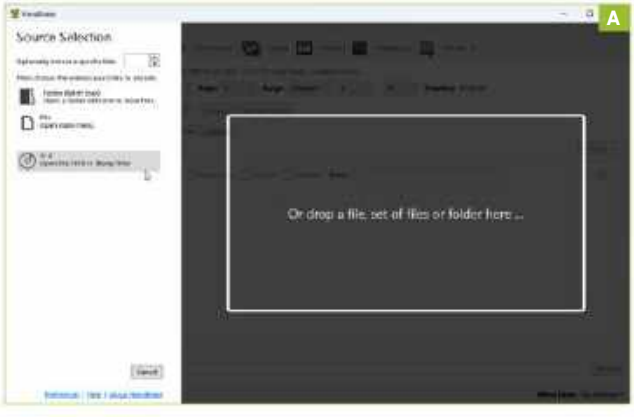

1 SET UP HANDBRAKE AND YOUR SOURCE FILE(S)

The best tool for converting files to a universal format is HandBrake. It supports all the standards defined in the box and enables you to craft the perfect stream containing your choice of audio and subtitle tracks. Start by downloading and installing the latest version from www.handbrake.fr.

» HandBrake can convert existing media files covering a wide range of formats and codecs. It can also process files from physical media—HandBrake can work with DVDs, ISO files, and some Blu-rays, or copy the files to your hard drive first using MakeMKV (www.makemkv.com).

» Once the files are on your hard drive, consider renaming them to make them more easily identifiable if required before opening HandBrake and clicking the 'Open Source' button. This opens the 'Source Selection' window [**Image A**], from where you choose to select a file or folder containing multiple titles using the buttons on the left, or drag files from a folder already open in File Explorer on to 'Or drop a file, set of files, or folder here.'

» Once selected, HandBrake opens and processes the files before placing you at its main screen, ready to prep the first file for conversion.

2 SELECT WHAT TO RIP

HandBrake displays a summary of the first matching title in the list, complete with preview image. Click the 'Title' drop-down menu to see what titles have been added if you selected a folder, ISO file, or disc. If you can't easily identify a file by its name, its length is displayed in the 'Title' field, or click 'Preview' to open a larger preview window with timeline slider. Failing that, open the file in a media player, such as VLC (www.videolan.org/vlc).

» To the right of the 'Title' field are 'Angle' and 'Range' fields. The 'Angle' option is only needed in those rare instances when the movie contains multiple video tracks—choose which one to embed in your new file—while the 'Range' field enables you to choose to convert part of the movie as opposed to all of it. You can select what to convert using chapters, seconds, or individual frames, depending on just how precise you want to be.

3 EXPLORE PRESETS

Beneath the 'Title' and 'Range' fields is a 'Presets' drop-down menu. Presets make HandBrake manageable on a day-to-day basis, and basically collect all the settings configured in the tabs beneath to enable you to easily switch between customized settings depending on what you're encoding—for example, you need different settings for 4K, HD, and SD movies, or to cover video with different audio streams (such as 5.1 surround versus regular stereo). You may even need to go further and have fine-tuned presets for different devices or for different types of content.

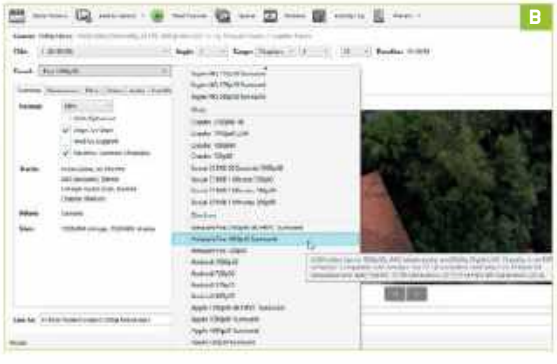

© HANDBRAKE

» HandBrake ships with dozens of presets, split into six categories: General, Web, Devices, Matroska, Hardware, and Production. Select one [**Image B**] and you'll see the 'Summary' screen change to reflect the settings being applied—explore the other tabs to see what other settings are applied.

» Each preset name is largely self-explanatory, with one key thing to decipher: At the end you'll see something like 1080p30 or 720p60. The 1080/720 refers to the vertical resolution, while the p30 or p60 refers to the number of frames per second. Typically, 25fps or 30fps is sufficient for broadcast material (25fps tends to be NTSC media, 30fps PAL)—the more fps and the higher the resolution, the larger the file.

» You should also pay special attention to the Hardware set of presets. You'll see many are grayed out. These highlight HandBrake's support for special x264 and x265 codecs that use your graphics chip's capabilities to speed up the encoding process. You'll see entries referring to QSV (Intel chips with onboard HD graphics), NVENC (Nvidia), VCN (AMD), and MF (ARM chips). Note, because these codecs have fewer tuning options, they're not generally recommended, because you invariably end up with larger files and potentially lower-quality video, too—only use them if your CPU struggles to convert files within reasonable timeframes.

4 CONFIGURE VIDEO OPTIONS

If you're in a hurry, one of the supplied presets will do the job, but for complete control over your files, we recommend either taking an existing preset and fine-tuning it or starting from scratch. Once you've constructed your own preset, click the 'Save New Preset' button to save it.

» Building your own preset involves exploring the seven tabs beneath the 'Preset' drop-down menu. The 'Summary' tab requires you to select which file container you plan to use for your file(s): MP4, MKV, or WebM. For maximum compatibility, choose MP4—make sure you check the 'Web Optimized' and 'Align A/V Start' boxes [**Image C**]. However, there may be times when MKV suits you better—specifically when converting HD content with subtitle tracks that are stored in a graphical format (PGS) rather than text (VOBSUB). Another potential advantage of MKV is that servers such as Plex and Jellyfin display the track titles of your audio and subtitle streams as opposed to just displaying the channels and/or language.

» Once your container is chosen, turn your attention to your movie's video stream, which is configured over the following three tabs. The 'Dimensions' tab is where you decide your movie's video resolution, aspect ratio, and control any required cropping—see the annotation (page 67) for details.

» Use the 'Filters' tab to try to correct problems with your source material. The Detelecine and Deinterlace filters are only needed with older broadcast material that produces a combing effect instantly visible in your preview window. The other filters help correct or minimize specific problems. For example, Denoise offers two filters to tackle excess grain or noise—NLMeans is the more effective but is slower. See https://handbrake.fr/docs/en/latest/technical/filters-summary.html for summaries of all filters.

» To see the effects of each filter, click 'Preview' followed by the 'Encode Preview' button to generate a short segment of your movie—30 seconds is the default, but you can encode up to 150 seconds if you need to.

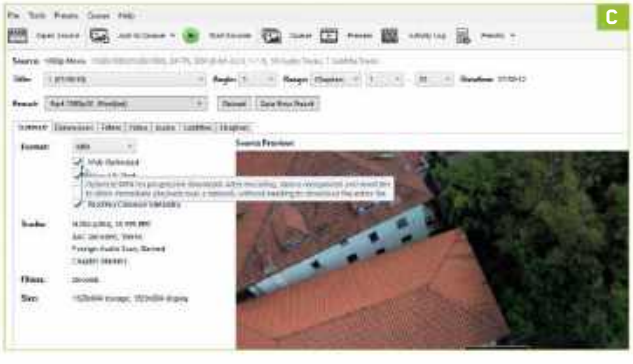

5 CHOOSE YOUR CODEC

The final tab for setting up video is arguably the most important. This is where you choose a video encoder and set options based on the movie you're encoding. Click the 'Video Encoder' drop-down menu, where you'll see a wide range of encoders are supported, including AV-1, MPEG-2, and MPEG-4, but for compatibility reasons, you should choose H.264 or H.265 (see box above to choose which is best for your movie).

» You'll see there are multiple H.264 or H.265 codecs available—in most cases, the default x264 or x265 CPU-based encoders work best. Choose 10-bit x265 when video quality (even for 8-bit sources) is more important than file size—10-bit files tend to be 25% larger than

PERFECT MEDIA SPEC

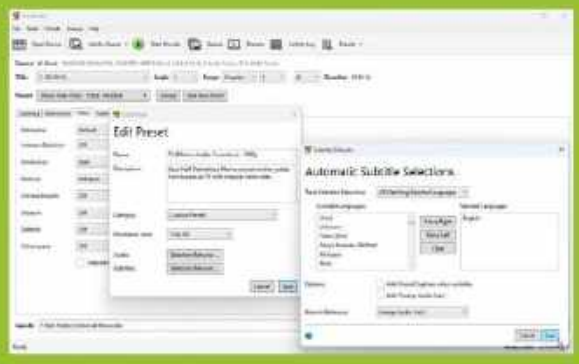

What does your universal media file look like? Let's start with the file's main element, the video codec. Here, you have a choice of two standards depending on your needs. For standard or high-definition content, the H.264 codec ensures your movies play on virtually any media player. If you're ripping UHD content—4K or 8K—choose the newer, albeit less universally supported, H.265 codec. Its main advantage over H.264 is that it produces smaller files at the same quality, keeping your UHD movie files a manageable size; however, it's not well supported on older players.

When it comes to audio, we recommend choosing the AAC codec for standard stereo tracks, and AC3 audio for all other tracks, including surround sound and—if applicable—commentary tracks. Because media files support multiple audio streams, why not play it safe and use one of each to cover all bases?

Your final choice is which file container to wrap them in—MP4 is more universally supported, but MKV is more versatile. Servers such as Plex can convert the file container to a compatible format without having to transcode the whole stream on the fly (a process called direct streaming), so it's not a deal-breaker choosing MKV.

TEST YOUR RIPS

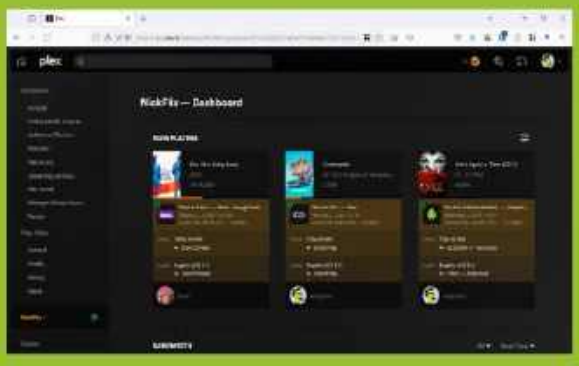

How do you know your converted files will work as expected? The simple answer is to add them to your Plex or Jellyfin library, then try streaming them from a variety of devices, both within your home network and remotely.

Once playback has started, navigate to app.plex.tv in your browser, then click the 'Activity' button, choosing 'Dashboard' from the drop-down menu. You should see the media being played appear under 'Now Playing', complete with information about its video and audio streams. The holy grail you're looking for here is 'Direct Play'—this means your server's doing nothing but provide the stream, and you should see that CPU usage is negligible.

If 'Direct Play' isn't supported, 'Direct Stream' is almost as good—the only change here is repackaging the file in a new compatible container with the client. The least desirable option is 'Transcoding', because the server must convert the entire file as it's being played. To prevent your CPU usage going through the roof, check to see if your server's graphics chip can be utilized to provide hardware-assisted transcoding—when enabled, you'll see '(hw)' appear next to 'Transcoding', and your CPU usage should drop as low as 10% rather than 90-100%.

their 8-bit equivalents. Beneath this, leave 'Framerate' set to 'Same as source' along with 'Variable Framerate'.

» Leave 'Constant Quality' selected, then place your mouse over the slider to reveal a pop-up with recommended settings depending on your video's resolution [Image D]. Move the slider left or right as required to adjust this figure in increments of 0.5.

» The 'Encoder Preset' slider works in a similar way to the 'Constant Quality' slider, except here you're trading off compression efficiency (so quality and file size) in return for a faster conversion time. The 'Encoder Tune' drop-down menu is a horses-for-courses option, so be prepared to change this depending on the type of movie you're converting, from 'Film' and 'Grain' to 'Animation'. Check the 'Fast Decode' box to aid playback on lower-powered devices. Finally, leave 'Encoder Profile' set to 'High' and set 'Encoder Level' to 4.1 for HD content, or 3.1 for SD content.

» Again, use the 'Preview' function to encode a small section of video to see the effects of each tweak without having to commit to a lengthy conversion of the entire file each time you want to make changes.

6 AUDIO SETUP

The 'Audio' tab is where you configure your movie's audio track—or tracks if you wish. Like movies on DVD and Blu-ray, video files can support multiple audio tracks, which the likes of Plex enable you to switch between when watching your movies. Each separate track increases the size of your file, but like the video track, audio tracks can be compressed to save space.

» One audio track is defined by default. The 'Source' reveals the currently selected track's number, its language, and its properties (the codec used, plus number of audio channels). Click this to see all available tracks, and to switch to a different one if you prefer. Alternately, click the 'Tracks' button to add additional tracks. These can be generated from the same source track, or you can choose a different track.

» The other columns define the track's properties after conversion. Start with 'Codec'—AAC (avcodec) is defined by default, which we recommend for two-channel stereo or mono tracks. For surround-sound tracks, we recommend AC3. Indeed, why not produce one of each track? This ensures maximum compatibility, plus gives you a choice of audio tracks that are optimized for both your bedroom TV and your den's home cinema.

» The remaining options determine the quality and number of channels in your new track. Quality is best left set to 'Bitrate', while 'Mixdown' should be set to the number of channels (such as two-channel Stereo or Surround 5.1) required. The other three options—'Samplerate', 'Gain', and 'DRC'—can be ignored.

» If the source track's codec is also AC3, then consider setting 'Codec' to 'AC3 Passthru' to avoid making any changes to it. One final step: Click the down arrow button to the right of 'DRC' to reveal the track's current name. By default, it's set to the source track's channels (such as Surround 7.1 or Dolby Surround), so you may wish to change this. It's also where you might provide details of a commentary track, for example.

© PLEX

SET YOUR OUTPUT RESOLUTION

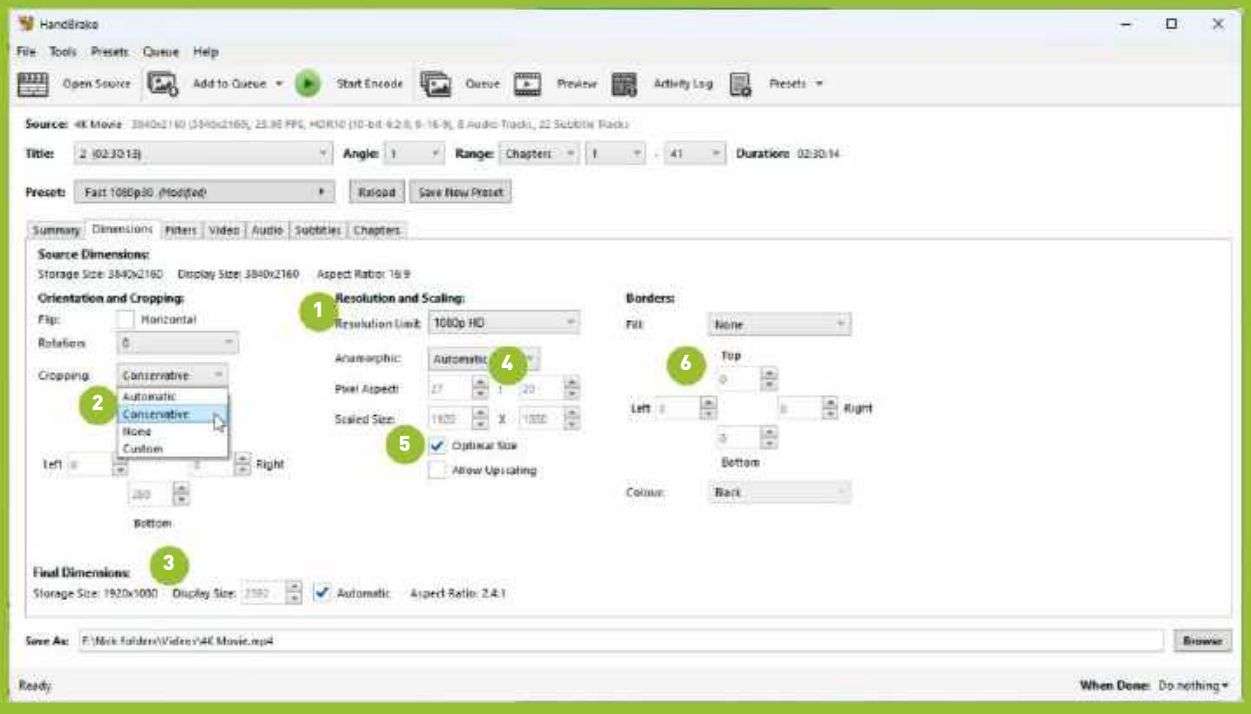

1. SET RESOLUTION LIMIT
Use this drop-down menu to quickly downscale your video to aid streaming over slower networks or to reduce the file size. Presets include 4K, 2K, HD (720p and 1080p), and SD.

2. CROP OUT BORDERS
Many source videos ship with superfluous black borders. Use these controls to automatically or manually trim them—use 'Preview' to see the effects of your changes.

3. ANAMORPHIC
Many videos on physical media are encoded in such a way that the video itself must be stretched or squashed to display correctly. This setting ensures that your video preserves that setting.

4. FINAL DIMENSIONS
If the anamorphic settings don't prevent the video looking stretched or squashed, uncheck the 'Automatic' box and adjust the 'Display Size' figure until the correct aspect ratio is displayed.

5. SCALED SIZE
Uncheck 'Optimal Size' if you'd like to manually specify the output resolution (using the 'Scaled Size' figures). Check 'Allow Upscaling' to increase the size of the video, accepting it may appear blockier as a result.

6. BORDERS
This does the opposite of cropping the image by enabling you to create your own borders—your video size is increased to accommodate them.

» Note: If you'd like to change the default settings for this tab, click 'Selection Behavior' to open a new dialog where you can configure the default track settings plus set up multiple tracks in one go [Image E]. Be sure to choose at least one language before setting 'Track Selection Behavior' to either 'First Matching Selected Language' or 'All Matching Selected Languages'.

7 ADD SUBTITLES

If your movie comes with built-in subtitles, you may want to include these in your final video. Switch to the 'Subtitles' tab, where you'll likely find a Foreign Audio Scan is already set up. This searches the movie for embedded subtitles where non-English dialog is spoken, but the tool is hit and miss, and adds unnecessary time to the scan.

» Click 'Selection Behavior' to set default actions for subtitles—it works in a similar way to the 'Audio' section, so choose which languages to include by default, and whether to include all instances or just the first subtitle track containing that language. You can also remove the 'Foreign Audio Scan' option from here, too. When done, click 'Save' followed by 'Reload'.

» If your movie contains embedded subtitles, they may either be already 'burned in' to the video track, so always visible, or they may reside in a dedicated subtitle track. To determine if this is the case, play the file in VLC and open the 'Subtitle > Sub track' menu. If one of the tracks is selected, this is likely your 'forced' subtitle track.

» Back in HandBrake, you have two options: You can make this subtitle track appear by default when the movie is played by checking 'Default', or you can burn it permanently into the movie by checking 'Burn In'.

» As with audio tracks, you can also name subtitle tracks to help identify them—this is useful with commentary tracks. Remember, the track name only appears in Plex and other media servers when you save the file in MKV rather than MP4 format.

» With your settings in place, you're ready to rip your movie—either click 'Start Encode' to process it immediately or 'Add to Queue' if you want to line up multiple files to process at once. And don't forget to save your carefully crafted setup as a preset!

Harness your VPN kill switch

YOU'LL NEED THIS

WINDOWS 11
An up-to-date installation.

VPN SUBSCRIPTION

USING A VPN (virtual private network) is a matter of common sense these days. For some internet users, routing their connection via a VPN server based outside their home country is also the best way to access popular websites. VPNs also offer excellent security, as all traffic between the device and server is encrypted.

Still if you're unaware the VPN connection has dropped, there's a risk you'll continue using the web in a way that lets others undermine your privacy. This is why some VPN clients deploy a kill switch, whereby all internet activity is halted until the secure connection is re-established.

Not all clients support this feature and those that do sometimes implement it in different ways. In this guide, you'll learn more about how to configure your kill switch (if available), as well as deploy some alternatives if your provider doesn't support it. **–NATE DRAKE**

1 CONFIGURE YOUR KILL SWITCH

The easiest way to get started with a VPN kill switch is to subscribe to a service that supports it. Some VPN providers offer a browser extension, but to use the kill switch, you need to make use of the client software.

» The specific steps to configure the kill switch vary by client but you can usually access them via the client settings. Some clients, such as NordVPN [**Image A**], enable you to fine-tune the kill switch; although its stated purpose is to prevent internet activity if your VPN connection drops, this doesn't usually apply if you manually disconnect.

» Check whether your chosen client allows you to also block the internet connection, even if you do so. This is the safest implementation of the kill switch, as bad actors with access to your device could sever your connection to the VPN server either by gaining access to your device or via malware.

» Even if your VPN provider claims to support kill switches, make sure the feature is enabled within your client.

» For instance, during our tests with the free tier of the Proton VPN client, the kill switch feature was disabled by default [**Image B**]. Users can enable the kill switch for specific countries or select 'Permanent' to enable it for all connections (recommended).

» Certain VPN clients enable you to make exceptions for local network connections, such as to a printer. These connections aren't typically made over the VPN interface, so can be allowed.

» If you're on an untrusted network, try to use a client that either disables local connections altogether or supports whitelisting only specific local addresses, such as the Windows 10/11 software offered by VPN provider hide.me.

» Configuring your kill switch in this way means that your device won't connect to the internet if the VPN client isn't running. In Windows, go to 'Settings > Apps > Startup' to make sure your client launches automatically upon login [**Image C**].

2 SET UP A MANUAL KILL SWITCH

If your chosen VPN provider doesn't offer a kill switch feature, technically you can manually configure Windows Defender to only allow outbound connections via the VPN.

» In practice, however, this was difficult to achieve in our tests, because even after entering the subnet mask of our VPN adapter into the firewall rules, we had to set up a separate rule blocking all other outbound connections that interfered with the VPN client. Although we encourage readers to play around

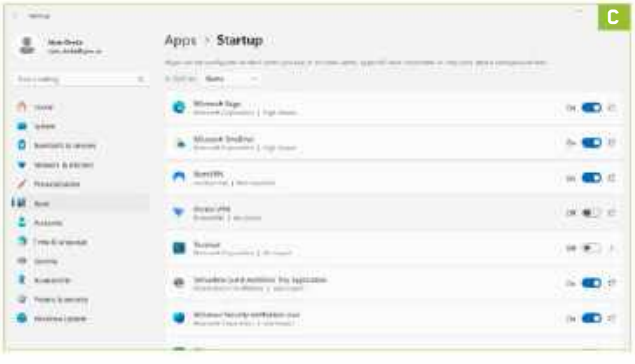

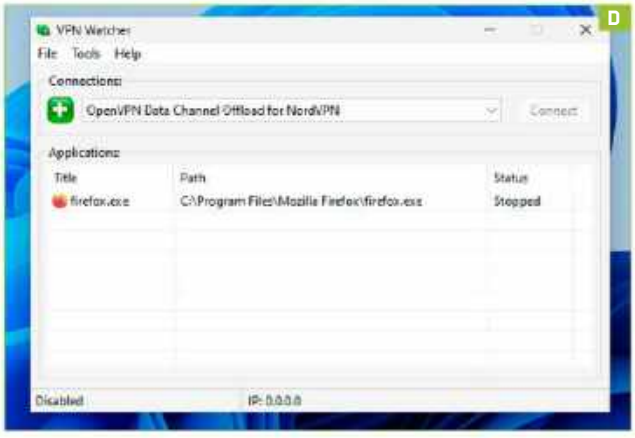

with their network configuration as they see fit, you may find it easier to use VPN Watcher instead (https://ugdsoft.com/products/vpnwatcher/).

» This program is capable of monitoring your VPN connection to see if it drops, then terminating the apps you specify. Note that the free version of VPN Watcher only supports doing this for one program.

» To get started, download and run the installer using the link above. Once VPN Watcher first launches, you'll see a simple list under 'Connections'. Certain services such as PureVPN and HMA are already supported, so you need only enter your credentials. However, network interfaces for other providers, such as NordVPN, should be available, too.

» Once you've selected your connection, make sure you close down any internet apps besides your VPN client itself. You can now click 'File > Add Application' to navigate to your chosen executable [**Image D**] .

» Once your chosen application(s) have been added, test your jury-rigged kill switch first by running one of them while the relevant VPN connection is active. Next, close or pause your VPN client app as well as your chosen program. Relaunch to make sure that it fails to run when your VPN connection isn't active.

» Paid versions of VPN Watcher enable you to suspend programs and run various command-line parameters. Still, the free version is sufficient if you simply want to disconnect your main web app when the VPN connection drops.

» During our tests of the Firefox browser, the kill switch kicked in even when we used a VPN browser extension, so make sure the client is active.

©VPN WATCHER

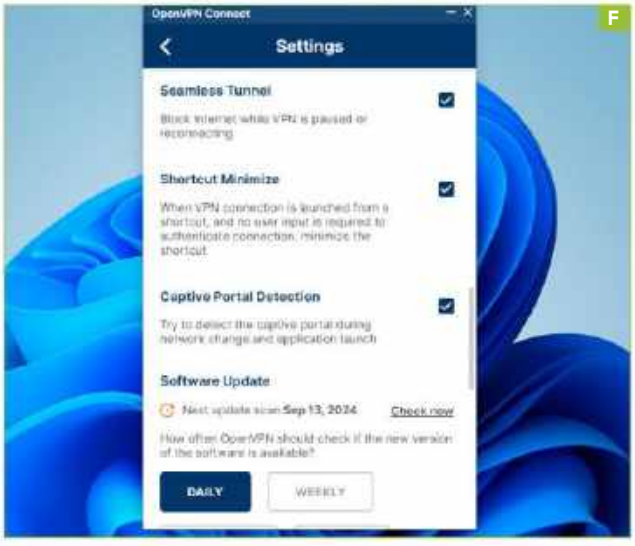

3 OPENING UP

If you want peace of mind when using your VPN even though the client doesn't have a kill switch, check if your provider supports the OpenVPN protocol.

» If the answer is yes, it's likely it will offer various OpenVPN (.ovpn) configuration files for download, which you can use to connect to the service via the OpenVPN Connect Windows client.

» This software is open source, so is much less likely to contain any undetected security bugs. Most importantly, it can block internet connections when the VPN connection isn't active.

» Begin by visiting https://openvpn.net/client/client-connect-vpn-for-windows/ to download and launch the setup assistant.

» Once the program launches, choose 'Upload File' then click 'Browse' to navigate to the .ovpn file you downloaded from your provider.

» Check the box marked 'Save Password' and enter the username and password supplied by your VPN provider. Note that the app-specific password may be different to the one you use to log in to your account [**Image E**]. If your provider also gave you an external certificate, you can specify it here, too.

» Once you're connected, click the hamburger icon at the top-left and choose 'Settings'. Scroll down to 'Launch Options'. From here, choose 'Connect Latest' to ensure OpenVPN Connects automatically launches on login and tries to connect using your latest profile.

» Next, scroll down to 'Seamless Tunnel' and enable the check box. As the description says, this blocks internet access while the client is paused or reconnecting [**Image F**].

» Settings are global, so apply to any other OpenVPN profiles you import. However, the 'Seamless Tunnel' isn't a true kill switch in that if you close the OpenVPN Connect client, other apps can still access the internet.

» Naturally, by connecting in this way, you'll also only be able to connect to those servers that support the OpenVPN protocol.

» If you're concerned about speed, ask your provider if any servers support DCO (Data Channel Offload). This implementation operates in the kernel space, so offers much faster speeds than traditional OpenVPN.

Write a novel with Quoll Writer

YOU'LL NEED THIS

QUOLL WRITER FOR WINDOWS, MACOS OR LINUX
(https://quollwriter.com)

DO YOU HAVE AN IDEA for a story rattling around in your head, but can't get it down on paper? The notion you can sit down in Microsoft Word and bash out a story from start to finish is enough to induce writer's block in even the most prolific wordsmith, while those who have diligently taken notes about character, plot, location, and other story elements struggle to pull them all together to deliver a coherent narrative.

This is where Quoll Writer steps in. It has everything you need to plan your story, write it without distractions, and even get feedback from other writers. And when you've finished, you can export it in Microsoft Word or ebook-friendly EPUB format to share with others.

In this tutorial, we'll introduce you to the key elements of using Quoll Writer to tell all your stories, including non-fiction works like family histories. You'll also find lots of handy hints in the program itself, as well as comprehensive online documentation, which you can access from the main menu or by going to https://quollwriter.com and clicking 'User Guide'. –NICK PEERS

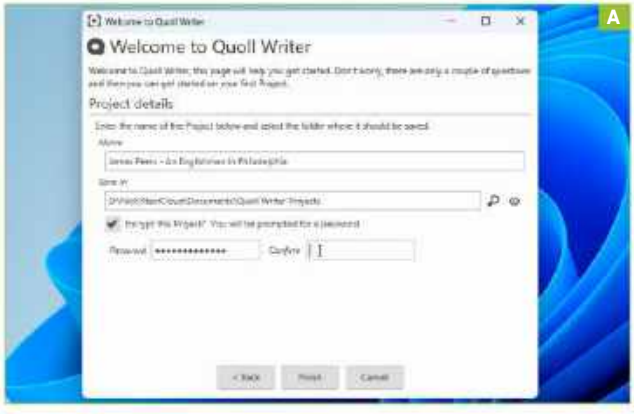

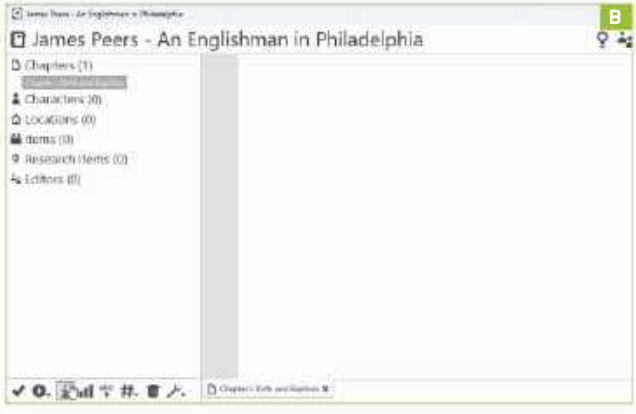

1 GET SET UP

Head over to https://quollwriter.com to download and install Quoll Writer. Once done, allow the program to launch. Leave English selected and click 'Next', read and accept the license, then choose 'This is my first time using Quoll Writer' before clicking 'Next'. Wait for the spellchecker to download, then click 'Next' again. Choose 'Create a Project' and click 'Next'. Give your story a suitably descriptive name—by default, it's saved inside the 'QuollWriter\projects' folder inside your user folder—click the magnifying glass to select a different location (such as one inside your 'Documents' folder). You're also given the option of encrypting your files for security reasons—check 'Encrypt this Project' [**Image A**] and provide a suitably strong, but memorable, password. Click 'Finish'.

2 BASIC NAVIGATION

Quoll Writer's main screen appears, which is split into two panes. On the left are shortcuts to your assets: 'Chapters' contain your story, while 'Characters', 'Locations', 'Items', and 'Research Items' will in time contain your project's notes and assets. At the bottom of this pane you'll see a series of icons enabling you to save your work, add new elements, access the current chapter information, view word counts and other statistics, toggle the spellcheck on or off, add tags, delete the current item, and access additional tools, such as the problem finder to help spot potential issues with your prose.

» Before typing anything, click the 'i' button [**Image B**], which will display the chapter's information sidebar in the left-hand pane. Here you'll see options for adding a description or summary for the chapter, as well as setting goals (such as how you want the chapter to progress the story). There's also space to record a plan, which could be a breakdown of what's needed to complete the chapter, and links, which in time will allow you to link chapters to assets and vice versa (these links can be clicked to jump directly to specific parts of your project). Once you're done, click the X at the top of the pane to return to the previous navigation pane.

© QUOLL WRITER

» Before moving on, click the cog icon in the top-right of the window to reveal Quoll Writer's main menu. Select 'Targets' from this list—the left-hand pane will display options to set yourself word count targets spread over sessions, days, weeks, and months. You can also set a limit on chapter lengths from here.

3 BREAK UP CHAPTERS

You can write chapters from scratch, or you might find it easier to break them down into scenes. Not only does this enable you to tackle your work in manageable stages, but you can also easily rearrange their running order later if necessary.

» The quickest way to add a scene is by right-clicking inside the main text editing pane on the right to choose 'New > Scene'. Provide a suitable description and click 'Save'. The description appears with a scene icon to its left, and you can start typing your story beneath this, or continue to break down the chapter further by adding additional scene titles [**Image C**]. You can also add other elements to the page: notes, plot outline items (rough outlines to move the story along without writing it in full), and 'needed notes' (a means to simply articulate what the story might need at this point).

» Note: You can swap the order of chapters by dragging and dropping them in the left-hand navigation pane; however, while you can drop in notes and plot outline items, you can't rearrange these. Instead, click an icon in the gray sidebar to reveal options for editing, linking, or deleting it.

4 GET WRITING

Once your story structure is sketched out, it's time to get writing. Quoll Writer also offers various aids to the writing process—the built-in spellchecker is one, but you can also tap into warm-up exercises (click the 'Settings' button to access this), as well as a 'Find Problems' tool. Once you're ready to dive into your story, select your chapter and click the full-screen button next to 'Settings'. Roll your mouse up to the top to reveal a hidden menu [**Image D**]—click the headphones for distraction-free mode. When you're done, press F11 to exit.

5 FINAL THOUGHTS

Quoll Writer also includes a stats tool with readability score, Idea Board for taking notes, and a free Editor's Service, where you can swap feedback with other users. Choose 'Options' from the main menu to customize the program further should you need to—focus on the 'How things look and sound' section to change the font type and size, plus rearrange the interface (click on the 'Interface layout' diagram to reveal a pop-up menu with five alternatives to choose from [**Image E**], plus choose where the toolbar and tabs are displayed), and configure dark mode to come on automatically.

» Once your story is ready for sharing, open the main menu and choose 'Export Project' to export it as a Word (.docx) or ebook-friendly (.epub) document. You can export all or part of your story, along with any research notes you might wish to include, which will be saved in a separate file from your completed work.

GET ORGANIZED

Keep your notes organized by attaching them to assets such as characters, locations, items, and research items. These work in a similar way—right-click one and choose 'Add a new...' to set one up from scratch. Fill in basic information—name and description—then click 'Save'. A new entry is created and displayed in the main pane on the right inside its own tab (see the bottom of the pane). Right-click to add new user-defined fields, which can contain all kinds of information, from text and images to web links and even attached documents. You can even arrange your fields across multiple columns to fit more on screen. Long story short, you can record just the basics or build up detailed information about any element of your narrative. Don't forget to create a dedicated 'Links' field to enable you to link one asset to another. Link assets to chapters via the 'Chapter information' sidebar, or click a 'Scene' as shown in the screenshot to link multiple assets to that scene.

» You can build detailed notes to support your story.

If none of the supplied asset types fit, you can even add your own (right-click inside the left-hand pane and choose 'Add New Type of Asset').

LAB NOTES

ZAK STOREY, CONTRIBUTOR

Hypervisor hell with AMD & NZXT

One little program that stole multiple days of my life

I'm a massive advocate of regularly re-installing Windows. Whenever anyone asks me for advice on buggy operating systems, it's my first port of call. Conflicts, bugs, corrupt Windows updates, odd registry keys, they all lead to issues, and most if not all can be fixed with a clean (internetless) install of Windows 11. I personally do it every three to six months.

This time, I decided to do a system swap, too, and move over to last issue's build (but with tweaks to storage and GPU; I was fed up with Intel and the 14th-gen drama). I swapped the hardware, got Windows installed, and began running through my program installs and updates. Everything was going fine, then bam. Hypervisor error, blue screen of death. Restart, log in, go again, and it was a constant feed of blue screens. I perused the internet for information but with little joy. Hypervisor errors can occur for all manner of reasons, from failed storage sectors, memory instability, and aggressive overclocks, to corrupt drivers. It's tied to the Hyper-V virtual machine systems on Windows itself.

After reinstalling Windows again, and getting the same errors, I decided it was the hardware. I swapped out the graphics card, SSD, and memory, and reinstalled Windows again. Exactly the same problems. Tried a third time with a third SSD, but something was causing this eternal loop of blue screens. I screamed into the abyss.

Eventually, I clocked that it was only after I installed NZXT CAM and configured the RGB lighting that things went wrong. A quick Google search later and I discovered that NZXT CAM, for its LCD displays on its AIOs, can piggyback off the iGPU in my AMD Ryzen chip instead of the dedicated GPU (saves resources), but for some reason this can cause a hypervisor error blue screen bug. NZXT has implemented a change to this—you can go into CAM's settings and ensure it uses the GPU not the CPU.

NZXT's Cam software, when it works, is flawless. When it doesn't, hold on to your hat.

One more Windows 11 reinstall later and I finally have a stable system, albeit without CAM running in the background. I'm destroyed. Man, this is not my month.

JEREMY LAIRD,

Contributor

MSI has rolled out its new X870 and X870E motherboards for AMD AM5 processors, and the big news is a boost in PCIe power supply. The boards' new supplemental PCIe Power feature comes courtesy of an integrated eight-pin PCIe power connector on the motherboard and delivers a major boost in overall board and GPU power. When you consider that next-gen GPUs are expected to be even more power-hungry, what with Nvidia's rumored 550W RTX 5090, that could be handy.

The connector also means you can deliver a bananas 132W just to the fan headers. You won't struggle to cool a PC based on one of these boards, no matter how batty you go with CPU and GPU specs.

But where does this end? The ATX standard seems to be falling ever further behind PC components and their power draws. Do we really want another connector? Wouldn't it be better if the 24-pin mobo power connector was overhauled so it could handle all power requirements?

Imagine one power cable for everything rather than the usual rat's nest inside a PC. It can't be an insurmountable engineering challenge, can it? As it is, we now have another supplementary motherboard connector for GPUs to go along with the existing one for the CPU. But you'll still have to plug yet another cable into the graphics card itself to keep it juiced up. Frankly, it's a mess and needs a total rethink.

Is native 4K and native 1080p actually possible in a single monitor?

Editor's Pick: Alienware AW2725QF

One monitor, two native resolutions

WHAT? ACTUALLY, the Alienware AW2725QF isn't unique. The LG UltraGear 32GS95UE ostensibly offers the same ability to run at both 4K and 1080p. But this Alienware does it differently. Where the LG's dual-mode capability is an interesting extra but the main attraction is the 32-inch WOLED panel, the Alienware monitor is very much defined by offering two 'native' resolutions.

That's because for the most part the Alienware AW2725QF is a very nice but not remarkable monitor. The basic specs involve a 27-inch 4K IPS panel that runs at up to 180Hz in UHD mode. It has DisplayHDR 600 certification, but no full-array dimming, so the HDR support is basic. But the response is rated at 0.5ms GTG, which is very zippy for an IPS monitor. Again, this is nice, just not radical. Except for that dual-native resolution thing.

The concept is straightforward: pixel doubling. You combine two pixels in both the vertical and horizontal directions for a grand total of four pixels from the panel's 4K grid acting as one to create a single pixel in an alternate 1080p mode.

That differs from the normal approach to rendering 1080p on a 4K panel, which sees the 1080p image interpolated to 4K using a simple spatial algorithm, the results of which look soft and blurry. Pixel doubling, aka integer doubling, means that each pixel should be much more precise.

What's more, as a dual-mode monitor, you can opt to have a higher refresh rate for the 1080p mode. Alienware has gone for 360Hz for the 1080p mode. The idea is that you have a single monitor that can do both the high-precision 4K thing for, say, single-player, strategy, and role-playing games, and 1080p high refresh for competitive online shooters. Nice.

If you're thinking, great, this alternate 1080p mode sounds like it should be just like an actual 1080p panel, it doesn't quite work like that. Without getting entirely into the weeds on the details, there are a few issues, including subpixel structure, pixel spacing, and pixel definition.

For starters, pixel doubling results in a different subpixel structure than a conventional single RGB pixel, which alters image quality. The shape of and spacing between pixels is different, too. For any given panel size, a true native 1080p pixel tends to be softer, rounder, and has more space between it and the next pixel.

Put all these elements together, and the result is that a pixel-doubled 4K panel running an alternate 'native' 1080p simply looks different to a panel of the same size that's actually 1080p. Is it better or worse? That's subjective. But like the LG OLED panel, you wouldn't mistake this for a conventional 1080p panel.

Compared to the LG, the results here are a little crisper and sharper. That's both a pro and con. It means this monitor avoids the slightly soft and blurry look of the LG. In game, it looks good much of the time. In fact, you'll often forget it's not truly a native 1080p panel. However, some graphical elements catch the pixel-doubling approach out. It's most obvious in fine details. A common example is grass and other foliage with fine lines. In 1080p, they can look horribly rough and pixelated.

It's worthwhile but not a killer feature. The rest of the monitor is very good, but would struggle to justify the price without the dual-mode shizzle. If you're not interested in the dual mode, there are better value 27-inch 4K options. **–JL**

$599, www.alienware.com

Reviewed...

Can the most expensive and powerful Zen 5 chip dominate the benchmarks and not just your wallet?

AMD Ryzen 9 9950X

AMD's top Zen 5 CPU looks very similar to its predecessor on paper

IF YOU DIDN'T KNOW anything about Zen 5, you'd think the new Ryzen 9 9950X is a step sideways from its predecessor, the Ryzen 9 7950X. They have the same number of cores and threads (16 and 32, respectively) and the same total amount of L3 cache (64MB). The 5.6GHz boost clock is the same for both, and the new chip even has a lower base clock (4.3 versus 4.5GHz). And yet, the Ryzen 9 9950X is definitely the better processor. The question is, by how much?

The answer is somewhere in the guts of the two CCDs (Core Complex Dies) that nestle next to the IOD (Input/Output Die) underneath the heatspreader. Indeed, the fact that the new 9950X has the same boost clock as the old 7950X means that any performance gains are going to be purely architectural.

Let's start with what the Ryzen 9 9950X is best at, and that's productivity, content creation, rendering, video editing, and so on. It's not just good at it, it's the best CPU for these tasks, period. Only a mega-expensive AMD Threadripper or Intel Xeon would likely achieve better results.

That said, it's only notably better in certain applications than the Ryzen 9 7950X. In the multithreaded Cinebench test, the 9950X is 15% faster than the 7950X, and 16% better in Blender. But the 9950X's lead over its predecessor is 8 percent in HandBrake and just 5 percent in our photo editing benchmark.

The performance caveats become more apparent in gaming. Compared to the Core i9 14900K, it generates 3 percent lower average frame rates, although the 1 percent lows are further behind, being 5 percent and 8 percent down respectively. That's close enough to not be noticeable in real-time gaming, and our other game tests put the Ryzen 9 9950X on par or slightly ahead of the Core i9 14900K. Against the old 7950X, the 9950X offers marginal gains, below 10 percent.

However, one aspect in all of this that strongly leans in AMD's favour is power consumption. Recording the average CPU package usage in *Baldur's Gate 3* and the multicore test in Cinebench provides clear evidence for this. Where the Ryzen 9 9950X consumed 129W and 196W in those two tests, the Core i9 14900K demanded 177W and 271W—that's 37 percent and 38 percent more power.

But what of the 9950X 105W Eco mode? In gaming, *Metro Exodus Enhanced* and *Total War: Warhammer 3* both took large hits to the 1 percent low figures with the 9950X limited to 105W. Other games weren't so badly impacted and the 9950X still outperformed all other tested CPUs in content creation apps with this limit.

All told, the Ryzen 9 9950X is an outstanding productivity CPU, but not the outright fastest when it comes to gaming. Arguably, the point is that it's more than good enough for gaming, while its multithreaded performance is second to none, and it's efficient, to boot.

The problem is twofold. First, none of this comes cheap. At $699, the 9950X has a $50 lower MSRP launch price than the previous model, the Ryzen 9 7950X, but you can bag the latter for a smidge over $500 right now. That makes the new Zen 5 chip around 35 percent more expensive. But it's not nearly 35 percent faster.

Which brings us to the second snag: It's not a huge step over the old 7950X. For sure, some applications will really favor the new architecture. But unless you're looking for a CPU for very specific workloads that happen to fly on Zen 5, you might as well save yourself over $170 and get the last-gen Ryzen 9 7950X. And if all you care about is PC gaming, avoid the Ryzen 9 9950X altogether, because there are better and cheaper options, such as the Ryzen 7 9700X and the mighty Ryzen 7 7800X3D. **–NICK EVANSON**

BENCHMARKS

	AMD Ryzen 9 9950X	AMD Ryzen 9 7950X	Intel Core i9 14900K
Cinebench 2024.1.0 (points)	**2,299**	1,993	2,044
Blender (samples/min)	**197**	170	149
HandBrake video encode (fps)	**132**	122	102
Photo editing (points)	**10,147**	9,683	9,151
Metro Exodus 1080p (fps)	**146**	145	141
Cyberpunk 2077 1080p (fps)	114	100	**118**
Average power consumption (rendering, W)	**196**	203	248

Best scores are in bold.

VERDICT 7

AMD Ryzen 9 9950X

WAX ON The fastest multithreaded CPU of all; Very efficient.

WAX OFF Not a big step over the 7950X; Very expensive.

$699, www.amd.com

SPECIFICATIONS

Cores	16
Threads	32
Base clock	4.3GHz
Boost clock	5.7GHz
L3 cache	64MB
L2 cache	12MB
Unlocked	Yes
Max PCIe lanes	24
Graphics	Radeon Graphics
Memory support (up to)	DDR5-5800
Processor base power	170W
Maximum package power	230W

Acer Nitro 14

AMD puts on a good show, but it still lacks punch

WE NEED TO stop with this AI nonsense. It's incredible enough what we can do with it, how it's radically overhauling medicine, research, and scientific endeavors. Why the heck do we need to advertise it in everything under the sun?

Take the Acer Nitro 14. This is a good gaming notebook range for the price. It starts out at $1,299, has a Ryzen 7 8845HS, 16GB of LPDDR5X (OK, it's soldered, that's not great, but you can get up to 32GB), comes with a 512GB SSD (specced up to 2TB), and up to an RTX 4060, all packed into a 14.5-inch form factor, complete with a 120Hz 1,920x1,200 IPS display (and in max config, again you can get that up to 2,560x1,600). What's the first thing we see on the product page? "Embark on an epic journey with the AI Acer Nitro 14 gaming laptop." Why? Just why? We all know that AI performance is solely dominated by the graphics card at this point. Yes, AMD's latest 8th-gen mobile chips do have AI in mind, and better performance thanks to an integrated NPU, but even in the best case scenario, even with the RTX 4060 in our review unit, with just a 100W TDP, they're going to be, what, 20x less effective? Of course, the marketing shifts to the fact that its TOPS performance predominantly comes from the GPU, but again, the card has been out since February 2023. Hardly revolutionary.

AI marketing faux pas aside, the Nitro 14 is an impressively capable notebook. In Cinebench 2024, multicore slides in at just shy of 800, and single core is just about into triple figures at 100 points, as well. The thing is, though, it's gaming that seems to be the Achilles heel of most modern gaming notebooks. And, yes, that does sound ridiculous. The sad fact is that screen resolution has come such a long way in such a short amount of time, that mobile graphics cards cannot keep up. They're thermally limited by the form factor, and as such, TDPs are often quite a bit lower as a result.

Take the RTX 4060 we have in our review sample. Complete with 8GB of VRAM, fortunately—unlike most of its mobile moniker siblings—like for like it actually matches its desktop counterpart for CUDA cores, Tensor cores, the works. Where it falls off, however, is the TDP, which has been reduced from 115W base to "up to" 100W, dragging the clock speed with it. What that means is that if you take this laptop into gaming, on its native resolution (our unit has the 2560x1600 screen), in something like *Cyberpunk 2077,* with DLSS and frame generation enabled, you're going to get an average frame rate of just 35. In *Total War: Warhammer 3,* that figure is 39fps. In something like *F1 2023,* you're looking at 25fps. That's not playable. The argument, of course, is that you shouldn't be gaming at this resolution, on this laptop, that you paid for, with its screen res at this state. No, that's solely for desktop use. Instead you should be gaming at 1080p, and to be fair, numbers do jump up quite considerably into the 60fps region and beyond if you do that. But it just makes it challenging to call this a gaming laptop. It's like buying a four-cylinder Ford Mustang instead of the same car with a V8. It looks great, but it doesn't run how you want when you put your foot down.

AMD's also still on the back foot when it comes to laptop efficiency. Although its 9000 series processors are really bringing the fight to the desktop realms, the same can't be said for its mobile range. Compare the Nitro 14 to the Predator Helios Neo 14, with its Core Ultra 7 155H, bigger screen, and same battery size. In modern office applications, the Neo 14 lasts 67 percent longer. Gaming is pretty much identical, admittedly, but that stat alone isn't enough to swing the fight in team red's favor.

The Nitro 14 is a good notebook. For the money, you are getting a solid proposition and a beautiful screen. It's not perfect, though, by any means. In fact, the gaming notebook market has inherent issues all the way across right now. Yet if you're looking for an entry-level(ish) laptop that can dabble in some gaming, while letting you get on with some work on the side, it's perfectly capable, at least as long as you keep it plugged in. Just ignore the AI nonsense. **–ZAK STOREY**

VERDICT 7

Acer Nitro 14

NO_2 INJECTION Solid all-around specs; beautiful screen; build quality is decent; battery life overall good; quiet.

LAUGHING GAS Can't really game at native resolution; soldered RAM; Intel CPUs still more efficient in day-to-day tasks.

$1,299, www.acer.com

BENCHMARKS

	Acer Nitro 14	Acer Predator Helios Neo 14	Alienware M16 R2
Cinebench 2024 Multicore (index)	799	**985**	979
Cinebench 2024 Single Core (index)	100	103	**105**
Cyberpunk 2077 Upscaling (avg fps)	75.4	62.1	**99.8**
Total War: Warhammer 3 (avg fps)	70.1	56.6	**93.5**
3D Mark: Timespy (index)	8,533	9,377	**12,116**
3D Mark: Timespy Extreme (index)	3,998	4,501	**5,667**
PC Mark: Modern Office Test (minutes)	233	**389**	363
Price (USD)	$1,299	$1,409	$2,049

All benchmarks performed with a clean install of Windows 11, with the latest updates, chipset, and drivers installed. Games tested at 1080p on Ultra preset. Modern office test performed at 60% brightness. Best scores in bold.

SPECIFICATIONS

CPU	AMD Ryzen 7 8845HS
GPU	Nvidia GeForce RTX 4060 8GB
RAM	16GB LPDDR5X @ 6,400MT/s
Storage	512GB M.2 PCIe 4.0 SSD
Screen	14.5-inch 1920x1200, IPS, 120Hz
Connectivity	2x USB 3.2 Type A, 1x USB 4 Type C, 1x USB 3.2 Type C, 1x 3.5mm analog out, 1x HDMI, 1x MicroSD slot, Wi-Fi 6E
Dimensions	12.76 x 10.07 x 0.8 inches
Weight	4.45lb

They say: "AI-powered laptop."
We say: "Nonsense."

The bezel is nice and thin, which adds to the elegant look.

Gigabyte FO27Q2

Speed is of the essence with this OLED screen

ONCE UPON A TIME, 60Hz was the standard. This weird hangover from NTSC TVs that needed to sync to the frequency of AC power (which in turn goes back to Nikola Tesla, Edison, and Westinghouse making arc lights) still provides a smooth viewing experience, but monitor manufacturers have realized they can go faster, and the rule of bigger numbers always being better is well and truly in play.

And while 240Hz may not be at the absolute cutting edge, it's still plenty fast, though you'll need a beefy GPU to push that kind of frame rate to it at 1440p.

This means it's kind of a niche product. There are loads of OLED screens that match the FO27Q2 for panel size (27-inch, in a standard widescreen shape) and brightness, and which may well cost less. The frame rate is something special that's going to appeal to lovers of competitive shooters and other eSports. An opponent can move a long way in 1/60th of a second, after all, so it's helpful to see them pop out from cover before you're killed.

There are other gaming features, too, such as the ability to add a crosshair to the center of the screen, and AMD FreeSync Pro is on by default to clamp down on tearing. There's more, in what Gigabyte calls its Tactical Features, such as Night Vision that lightens shadows without blowing out the whole picture, and the ability to monitor hardware stats in an overlay. This is a gaming screen through and through, and it's less well suited to creative work, something underlined by a color response that displays all of the sRGB gamut, but only 75% of the more pro-focused Adobe RGB and P3 color spaces in our tests. The manufacturer claims it can do better, so you may have better luck by adjusting the settings.

Still, if you buy it as a responsive gaming screen, you're not going to be disappointed. There are four video inputs at the back, including a useful USB-C KVM that can supply 18W of power to your laptop, allowing an elegant single-cable connection assuming you have a machine with the right kind of USB port. As a standard widescreen monitor, plugging a streaming stick into one of the HDMI 2.1 ports gives you something approaching a viable video-watching solution, though the pair of 5W speakers should probably be supplemented with something plugged into the 3.5mm audio socket. They're fine for Windows beeps, but game and movie soundtracks deserve something larger.

Build quality is up to the high standards we expect from Gigabyte, with a slender, graceful foot connecting to the well-adjustable stand. It has a widely splayed support that's big enough to keep a very small computer or speaker between, and the whole thing looks very nice on a desk alongside a small-form-factor PC.

For a 27-inch screen, 1440p is about right. Sure, you can squeeze 4K into the same sort of space, and 1080p screens of this size are plentiful, but by keeping pixel density in check, its sharp picture can still be driven at exceptional speed, while OLED's low response rate and excellent brightness and contrast keep ghosting at bay. The Gigabyte FO27Q2 is an excellent upgrade if you're playing *Counter-Strike 2* on an old HD LCD, but if you need something more versatile and cheaper, maybe look elsewhere. **–IAN EVENDEN**

VERDICT 7

Gigabyte FO27Q2

SCREEN A fast and sharp OLED monitor with plenty of gamer-centric features.

SCREAM A tightly focused product that may not be attractive to non-gamers.

$700, www.gigabyte.com

SPECIFICATIONS

Screen size	27-inch
Screen type	OLED
Resolution	2560x1440
Max refresh rate	240Hz
Color gamut (stated)	99% P3
Inputs	2x HDMI 2.1, 1x DisplayPort 1.4, 1x USB Type-C, USB hub, 3.5mm audio
HDR	VESA DisplayHDR True Black 400
Webcam	No
Features	KVM
Speakers	2x 5W
Adjustments	From -5°to +20° tilt, ±20° swivel, -0°+90° pivot, 130mm height
Dimensions	24x14x2.3 inches
Weight	15 lbs

AndaSeat Kaiser 4 XL

A chair as big and as comfortable as they come

GAMING CHAIRS come in all shapes and sizes, but if you're of larger dimensions, your options narrow down a bit. If you've got wide shoulders, are particularly tall, or simply need a bit more room, you need something like the AndaSeat Kaiser 4 XL. This substantial chair has serious dimensions—the backrest measures 34 inches tall, with a shoulder width of 22.5 inches, while the bottom cushion is 21 inches across and 20 inches deep.

This means that, even more so than standard sized chairs, this requires a friend to help put together, as the extra dimensions make for more of a workout than most. Once constructed, though, it's large rather than excessive— it'll fit most people without being too big to fit in with your average desk setup. It's a looker, too, with bridge-of-the-*Enterprise* vibes, set off by those wide, swooping shoulder wings and plush side supports.

Sitting in the Kaiser 4 XL for the first time is a little like getting on a plane and realizing you've accidentally been upgraded—there's much more room here than you would expect, and it's a pleasant surprise. This reviewer is not the tallest or largest of human beings, but it doesn't feel like he's surrounded by too much chair when he sits down. His 5' 11" frame fits nicely, while his 5' 4" partner doesn't look too out of place, either.

There's a ton of adjustability settings here, all of which work well. There's 135° of recline, accessible through a sizeable right-mounted handle (top marks for not putting this crucial control underneath the chair, so you can actually adjust it while leaning), and a pop-out integrated lumbar cushion, which can be adjusted by a dial on the side of the backrest and popped in and out with a lever on the side of the seat.

We'd go as far as to say the lumbar support here is actually the best we've used to date. It's easy to dial in exactly as much or as little as you need, and the pop-out cushion treads the line between being soft enough for comfort and firm enough to provide proper, meaningful back support.

There's also a magnetic head cushion with some seriously powerful magnets. It's good fun to chuck it at the chair and watch it attach with a pleasing thunk, and it actually requires some genuine effort to move it around once attached. Then there's the armrests, which are 5D. For the uninitiated, that means they adjust up and down, forward and backward, left and right, and rotate. Oh, and they split in the middle to fold upward, too.

If we were ordering our own model, we we'd pick the fabric version, because we find PVC leather chairs uncomfortably sticky in summer and cold in winter. We're not quite sure how AndaSeat has managed it, but the faux-leather model we received suffers from neither. On hot days it's been cool and refreshing, without causing excess sweat, whereas on cold nights, it's not frozen solid like the rest of our house.

While the materials feel premium, the design sometimes veers too heavily into gamer territory, with the rear very loudly displaying both "AndaSeat" and "Kaiser" in large, embroidered lettering. It's a handsome object overall, so a more subtle, smaller text design would have tied in better with the overall vibe.

For downsides though, that's pretty much it. It's well priced at $539, given our top gaming chair recommendation, the Secretlab Titan Evo, is $10 more expensive. You get a whole lot more chair (in terms of size, at least) for the cash here, making it a good value proposition.

Big, supportive, and very comfortable, the AndaSeat Kaiser 4 XL has room for all, with excellent lumbar support and plenty of adjustability. **–ANDY EDSER**

VERDICT **9**

AndaSeat Kaiser 3 XL

EMPEROR Good looks; roomy; pop-out lumbar support is very comfy; magnetic head cushion.

NEW CLOTHES Branding isn't subtle; soft backing prone to damage; armrests have slightly too much wobble

$539, www.andaseat.com

SPECIFICATIONS

Rec. height	5' 11" (181cm) to 6' 10" (210cm)
Rec. weight	176-397 lbs (80-180kg)
Recline	135°
Features	Magnetic memory foam head pillow, four-level pop-out lumbar support
Material	PVC leather (reviewed), linen fabric
Armrests	5D, upward folding
Colors	Black, white, purple, orange, brown, maroon, pink, blue (linen—gray and black only)

This big and wide chair is like getting an upgrade on a flight.

Gigabyte MO34WQC

An ultrawide that won't break the bank—or your desk

WHEN IT COMES TO ultrawide monitors, a 49-incher, although impressive, can be too big. The comparatively humble 34-inch screen is a great way to get the curved ultrawide experience without having to buy a bigger desk or watch YouTube videos on how to remodel a wall.

This OLED model from Gigabyte uses the 21:9 aspect ratio, which is like having one and a half 16:9 screens melded together after a teleporter accident, and the 1440p vertical resolution means there's plenty of screen real-estate. It's speedy, too, with a maximum refresh rate of 175Hz and AMD FreeSync Premium to line those frames up.

At the back, you get a decent array of inputs, including USB-C with a USB hub through which you can quickly share keyboard, mouse, and maybe a flash drive between two computers. It's super-handy if you've got a pair of PCs, and you can get 18W of charging for your laptop out of it, too, as long as you've got it connected to the right kind of port (video, data, and power, together at last). The ability to use the inputs in picture-in-picture and side-by-side mode is there, too, all baked into a complex but easy-to-use OSD menu that comes with the handy ability to customize the length of time it stays on the screen, which is incredibly useful if you're trying to set it up to get the best picture.

Build quality is excellent, with the stand neatly clipping into a recessed mount on the back, which also functions as a VESA mount if you want to hang it from an adjustable arm. The foot attaches to this stand with a single screw, and the whole support structure is very robust. You don't get a huge mount of adjustment—no rotation, 25° of tilt, and around five inches of height change—but what's there is probably enough to avoid having to balance it on a dictionary.

Performance from this kind of OLED panel is always good, and while the brightness of 250 nits in SDR mode (we measured it at 230 nits) isn't the best, pump in some 10-bit HDR content and you see a much better result, with Gigabyte claiming 1,000 nits over 3% of the screen, and VESA DisplayHDR True Black 400 for the rest of the time. You also get the usual OLED contrast and color saturation, and the 1500R curve, which is fairly gentle, and brings the edges into your peripheral vision if you sit in the right spot. It suffers from common OLED drawbacks, such as reflectivity despite an anti-reflective coating, and the way it picks up smudges.

A 34-inch screen like this is great for gaming—thanks to its resolution and refresh rate—and productivity. The ability to tile multiple documents side by side with no dividing line, or push the interface of a video-editing app into the periphery while keeping the content in the center, can be very useful. If you want to upgrade from a pair of HD 16:9 screens that always seem to be misaligned, or just want something elegant with the ability to switch between a desktop PC and laptop, this is a good choice. But there is a lot of competition in its price bracket, so it might look even more tempting if you can find it at a discount. **–IAN EVENDEN**

VERDICT **8**

Gigabyte MO34WQC

+ MONITOR Compact ultrawide that offers good brightness and refresh rate.

– MERRIMACK Bigger screens are available, and it gets smeary easily.

$949, www.gigabyte.com

SPECIFICATIONS

Screen size	34-inch (21:9)
Screen type	OLED
Curvature	1500R
Resolution	3440x1440
Refresh rate	175Hz
Color gamut (stated)	99% P3
Inputs	1x DisplayPort 1.4, 2x HDMI 2.1, USB Type-C (plus hub), audio
HDR	VESA DisplayHDR True Black 400
Webcam	No
Features	AMD FreeSync Premium
Speakers	2x 3W
Adjustments	Tilt -5 to +20°, height 130mm
Dimensions	32 x 23.4 x 9.6 inches
Weight	21.7lb

With a DisplayPort, two HDMIs, and a USB-C input, you'll be able to use it for several devices.

Acer Predator GM712

2008 called, it wants its projector back

IN A WORLD FILLED with BenQ and XGIMI projectors, it's becoming increasingly difficult to compete without bringing something either incredibly unique or remarkably well built to the table. Sadly, Acer's Predator GM712 does neither. And Acer has also added a healthy dose of model number confusion to the mix with this one, which makes it less than appetizing for us reviewers.

So, what's the beef? Well simply put, there are actually two GM712s, the model we have here on review, recently updated with VRR support, and the older version that, well, lacks it, and that's about it. The problem is, they both have the same product name. They're both listed as GM712. To tell the difference between them, you need to look at the full model name in the manufacturer notes. In the case of the VRR unit, you're looking for MR.JUX11.002 as opposed to MR.JUX11.001. Yep....

Another, perhaps more noticeable, difference between the two is the pricing. Right now, both are still in circulation depending on where you shop, and the 001 comes in at just under $1,000. Now that makes it a tempting opportunity (for a 4K gaming projector, with impressively crisp picture quality, and decent color accuracy). Particularly if it supports variable refresh rate. However, it's only the 002 version that has that latest feature, and as a result, the price bumps up by a substantial $300 for the privilege. Yeah, that's less than ideal.

CORE VALUES

Still, it's not all bad news. The Predator GM712 is a fantastic projection unit at its core. It's got full support for 4K 60Hz, 1440p 120Hz, and 1080p 240Hz. Contrast ratios are up there at 10,000:1, and its bulb life isn't bad either, with it rated at 15,000 hours on eco mode. Connectivity, similarly, is something of a dream. There are twin HDMI ports, VGA in, USB out, RS232, SPDIF optical out, analog out, and a wireless dongle that gets you 802.11ac Wi-Fi, as well as Bluetooth 5.0 support.

Unfortunately, there isn't any Google TV integration or anything along those lines, as you would find in the BenQ or XGIMI solutions, but given the price of the 001 unit, that's not the end of the world. Particularly as you're getting that image quality for a little under $1,000 (without VRR).

Using it, however, is a fantastic experience. We have a specific test here at *Maximum PC* where we shine every projector we get up against a gray office wall. This does two things. Firstly, it enables us to see how it handles projection on different wall types, and secondly, it demonstrates how intense the brightness can go. The GM712 performs admirably. Brightness is punchy and color rich, despite the gray wall. It performed far better than the murkier GD711 we tested earlier this year.

The problem is that it's not capable enough compared to some of the other options on the market. BenQ's X500i, for instance, runs rings around it when it comes to projection clarity and color pop, and comes in at not a lot more the 002 version, at around $1,500. Combine that with the fact that Google TV is integrated as standard (albeit with an added dongle) plus it has better audio, and it's a more appetizing proposition overall.

Adding to the GM712's woes, its availability also seems to be quite sparse, dropping in and out of stock in North America sporadically. What you're left with is a confusing product, with two slightly different specs, at massively different prices, that's often outperformed by the competition. At its core, it's a solid 4K gaming projector, and without the VRR support (which isn't the be-all and end-all), it's a tempting pick at under $1,000. Entering the $1,300 market, though, it's a whole other ball game, and you're better off looking elsewhere. –ZAK STOREY

VERDICT 7

Acer Predator GM712

ASTRAL PROJECTION Solid 4K performance; Decent latencies; Good pricing for 001 model.

PROJECTING INSECURITIES Outdated chassis; Confusing pricing and models; No Google TV integration.

$999 or $1,299, www.acer.com

SPECIFICATIONS

Resolutions	4K60, 1440p120, 1080p240
Display colors	1.07 billion (30-bit)
Contrast ratio	10,000:1
Brightness	4000 lumens
Light source life	15,000 hours
Max screen size	300-inch
Zoom ratio	2x
Projection offset	116%
Connectivity	2x HDMI, 1x VGA, 1x USB Type A, 1x RS232, SPDIF optical out, 3.5mm analog out, 3.5mm analog in, wireless dongle (2.4/5GHz), Bluetooth 5.0, Google Cast, Apple AirPlay

Razer Wolverine V3 Pro

Razer's pro controller is even better than Microsoft's

RAZER'S NEW WOLVERINE controller combines its excellent mouse switch technology, Hall Effect sticks, and an ergonomic design that we would argue is better even than Microsoft's Xbox Elite Wireless Controller Series 2.

Given the $200 price point here, you'd expect an impressive package, and that's what you get. It comes with a wireless dongle that works with both Windows PCs and Xbox Series consoles, a hefty carry case, a 10ft braided USB cable long enough for use in a living room, and two additional sets of thumbsticks (one concave, one convex). It's a package that includes everything you need, as well as being highly customizable.

The controller itself delivers excellent first impressions It feels chunky and solid; sturdy, like a quality product. The ergonomics are completely on point, at least for our not-so-massive hands. The controller just sinks into your grip, and the rear paddles are sublime. Razer has also added additional bumper buttons on top, which can be programmed using the app.

So, it's built well, expertly designed, and feels as though it'll last the course. It isn't trimmed in a soft-touch finish, but that doesn't bother us, as it's one less thing to get hammered over time. Perhaps more surprising is the lack of RGB, with just the company logo on the face of the controller lighting up.

When using the pad for the first time, the most noticeable feature is the clickiness of the buttons, rear paddles, triggers, and additional bumper buttons, as they use Razer mouse switches. That'll feel weird if you're coming from most other controllers, which are much softer when registering a button press.

Why bother? Performance, of course. Razer's mouse switches are designed for lightning-fast response and actuation, and that's exactly what you get here. It's hard to actually measure, but the fact it's here at all means you're getting the best of the best in this controller. Razer's gaming mice are legendary, and that same technology has now jumped to a controller.

The Wolverine V3 Pro utilizes Razer's HyperSpeed wireless, rather than Bluetooth or Xbox Wireless, and as such, you need to surrender a USB port on your PC or console to connect. In our testing so far, the wireless connection has been solid as a rock, and all appears fast and responsive.

ASSAULT ON BATTERY

The only downside of this, especially compared to a Microsoft controller, is the battery life. HyperSpeed wireless will simply drain the battery faster—that's just how it is. We wouldn't call it horrible at 12 hours, but you will be charging it up more often. Of course, you also have the option to hook up the included 10-foot-long USB-C cable to charge while you play or just to play wired all the time and forget about battery life.

Wired play really comes into its own on the PC. When you hook it up to your Windows 11 PC, you unlock 1,000Hz polling, something that makes the Wolverine V3 Pro stand out from the competitive crowd. You need to enable it first in the companion app, but the option's there, so why not use it?

The other big new addition to be excited about is Hall Effect sticks. You want these because they're just better. They're more accurate and responsive, and—most importantly—boast a much longer lifespan than traditional analog sticks, which tend to suffer from drift over time.

Whether on console or PC, Razer has a dedicated app to set up its controllers. There's no need to use Razer Synapse on the PC, and while this app won't win many awards, it's does what it needs to do.

There's not much competition in the high-end controller space, and Microsoft's Elite Series 2 has more customization options and Bluetooth support. However, the clickiness of the Razer's buttons offers a different feel, one that's more precise and satisfying to use. –RICHARD DEVINE

VERDICT 9

Razer Wolverine V3 Pro

ADAMANTIUM Excellent wireless performance; Hall Effect sticks; mouse switch tech; 1,000Hz polling rate, carry case.

STAINLESS STEEL Middling battery life; requires a USB port; no Bluetooth.

$200, www.razer.com

SPECIFICATIONS

Compatibility	Xbox Series X, Xbox Series S, PC
Connection type	Wireless (2.4GHz dongle), wired (USB Type-C)
Battery life	12-13 hours
Features	Hall Effect thumbsticks, circular d-pad, microswitch buttons, two-way trigger locks, Razer mouse click paddles, 1,000Hz polling rate (via wired connection on PC)
Accessories	Carrying case
Software	Razer Controller app (Xbox and PC)
Warranty	One year

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

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 $39.99 (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.

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

Acronis True Image 2025 Advanced for $29.99

ACRONIS TRUE IMAGE 2025 combines both a full backup suite and comprehensive anti-malware protection. It offers the full gamut of backup options, from individual files to complete drive imaging, and active disk cloning for migrating to faster or bigger hard drives.

The Advanced package improves on Essentials by offering you 50GB of cloud storage to back up your most important files and data remotely. You can still back up to local disks, network attached storage, or remote destinations via FTP. You can also back up Microsoft 365 data and set up automatic replication. It includes a one-device license for either a PC or Mac with a one-year subscription.

Remember to use code PRONOV20 at checkout for 20 percent off.

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

software.maximumpc.com

Razer BlackWidow V4 75%

Solid and compact, with hot-swappable switches

IF YOU'RE DECIDING on the right size of gaming keyboard for you, Razer may have an answer. The Razer BlackWidow V4 75% aims to pack in all the gaming prowess of its larger sibling but into a much smaller footprint. This newer iteration comes in smaller than tenkeyless, but a touch taller and wider than 60%. It gets a lot right for the most part, even if there's little that's exciting about this revision.

The main sell of the Razer BlackWidow V4 75% is that it's fully hot-swappable. Priced at $180, it's more affordable than the competition offering this feature. The Asus ROG Azoth 75% is $250 and the full-size Mountain Everest Max is $200. It's rare that the Boomslang brand comes out as the king of value, given its boutique status and branding, but hey, we'll take it.

Our review unit comes pre-loaded with third-generation Razer Orange tactile switches, which are said to deliver a quiet typing experience. We can verify this; they're notably quieter than the manufacturer's Green Clicky switches but a touch louder than its Analog Optical switches that you'll find in the likes of the Razer Huntsman V3 Pro TKL. While no other sets are included, you can buy new switches from Razer for $25 a piece.

Handily, Razer has included a multi-tool in the box for pulling keycaps and switches. The Razer BlackWidow V4 is held together with just five Phillips head screws for ease of access. Once you've lifted the lid, you'll see the tape-enhanced PCB, the plate foam, and the plate itself, into which you can slot the new switches. The gaming keyboard supports both three and five-pin options, so it takes the manual labor of having to build a board entirely from scratch.

Razer claims that the BlackWidow V4 75% has been made with an "optimized typing experience" in mind. As there's no way to change the actuation, the company has fine-tuned a solution instead. From our experience, we can say that it's among the best gaming keyboards that we've typed on, which surprised us, given the form factor. It feels like a bigger deck than it actually is, despite taking up less room than a TKL equivalent.

What's neat is how the design ethos of the full-size model has been carried over to this 75% version. There's a shrunken-down media roller in the top-right matching the original in all but length, plus the media buttons. You now get two instead of four, acting as pause and mute, but they get the job done. What's won us over, though, is the extra-plush wrist rest, which is incredibly comfortable.

Speaking of the form factor, the 75% variant does a solid job of skirting the line between size and function. It's a touch smaller than a tenkeyless model, but larger than a 60% variant as mentioned above. That means you're benefiting from a full function key row and page keys.

Gaming on the Razer BlackWidow V4 75% is a great experience. It feels sharp and responsive in story-driven games like *Assassin's Creed,* as well as in more immediate titles such as *The Finals.* Just a couple of games in, we were used to the smaller travel time of our hands to hit the keys against our usual full-size deck.

Of course, it wouldn't be a Razer product if it didn't feature Chroma RGB in some capacity, and this gaming keyboard is no exception to the rule. There's per-key lighting that can be tailored to your liking in the Synapse software, and the lighting is bright and vivid, even in the daylight hours, and looks especially good contrasting against the black design.

It's commendable that Razer was able to make this for $180 without sacrificing on build quality, so it's well worth a go if this is your first foray into keyboard modding. –ALEKSHA MCLOUGHLIN

VERDICT 7

Razer BlackWidow V4 75%

HOT Easily swappable switches; no compromises; rugged build quality; plush wrist rest; good pricing for functionality.

SWAP Ultimately unexciting; more expensive than full-size model.

$180, www.razer.com

SPECIFICATIONS

Size	75%
Base	Hot-swappable
Switches	Razer Orange Tactile
Backlighting	Per key
Anti-ghosting	N-key rollover
Polling rate	Up to 8,000Hz
Weight	1.8lb (815g)
Connection	USB-C

Elgato Wave Neo

The Kanye West of mics: Easy to get into, but a little weird

THE ELGATO WAVE Neo is the newest mic from the streaming gear company. Following a legacy of quality mics such as the beloved Wave 3 and Wave DX, the Neo line-up is aiming to deliver more affordable solutions. Thus, the Wave Neo costs $90, making it the cheapest Wave microphone to date. But it does make some sacrifices to hit that price point. For starters, the thoughtful recycled cardboard box contains only what you need. You get a braided USB Type-C (device side) to Type-A cable, a mic stand complete with some cable management, and the Wave 3 itself. The mic is rocking a pop filter natively.

The stand is quite good, but admittedly a little weird. It has a solid circular base with an 8cm diameter and standard quarter-inch screw connector accessible underneath. The first half of the pole screws in simply, and has a cable grip at the back to keep things in line. The second half is attached to the mic via another smaller screw connector, offset to the side. It slides straight into the first pole using a solid inner metal piece for strength and a tabbed connector. If it sounds confusing, it can actually only go together one way, so it's pretty simple.

The stand brings the mic to a height that's perfect for talking, but doesn't lock in a way that stops it from being lifted out. You may learn this the hard way if, like us, you pick up your mic and have that solid base fall right on your foot. Ouch.

The Wave Neo mic can be pivoted backward and forward at various angles and will hold that position nicely. On the back you have an aux port for monitoring your audio and the recessed USB Type-C port. The front sports one of the largest and brightest capacitive touch mute LEDs we've ever seen. Mercifully, it can be dimmed in software. But take care—the touch area is large and super-sensitive, and it's all too easy to unknowingly mute the mic.

Indeed, even a brush with a braided cable can set it off. The sheer size of the mute sensor is the main reason we wouldn't really recommend this mic to streamers or content creators. Finding out you've accidentally muted yourself mid-stream and no one noticed, or just didn't bother to tell you, is the worst.

WARM-HEARTED

Still, for the price, the voice recording is quite good. The sound isn't quite as crisp and punchy as something like the Wave 3, but it's warm and clear, which is what many will likely want for streams, gaming sessions, or meetings. However, with no controls other than mute on the mic itself, if you need anything dialed in or want to make changes, you're out of luck without introducing some software.

The Neo works with the Wave Link software, which is excellent. Because our test room has a lot of low computer fan noise and these are reasonably sensitive mics, it can be a huge problem. But we were able to easily turn the gain down enough to not pick that up. Of course, you've also got plugins like Elgato's own noise removal, if you need more help there, and as a plus it lets you keep the gain up for louder recordings. This all works well with Elgato's Stream Deck, too, which can help mitigate some of the lack of controls on the Neo itself.

All in, Elgato's Wave Neo is a solid-sounding mic for the price, and will bring a touch of warmth to your content. But in terms of recording quality, it's not kicking out our top budget or even top mid-range mic choices any time soon. The lack of on-device controls can be limiting, too. But if the convenience of the Neo has you sold, just make sure to leave that mute LED nice and blinding. –HOPE CORRIGAN

© ELGATO

VERDICT 7

Elgato Wave Neo

+ WAVE TO THE CROWD Cheaper than previous Wave offerings; Warm, crisp sound quality.

– UNDER THE WAVES Limited on-device controls; Mute button is triggered all too easily.

$90, www.elgato.com

SPECIFICATIONS

Condensers	1
Directional patterns	Cardioid polar pattern
Frequency range	20-20,000Hz
Recording sample rate	24-bit
Controls	Mute

Elgato Facecam Neo

Always sees us in the best light but lacks focus: 8/10 boyfriend, er, camera

THE ELGATO FACECAM MK2 is one of our favorite webcams. But the new Elgato Facecam Neo retails for $50 less than its predecessor at just $99. This makes it a conspicuously budget-friendly contender.

It comes with a matching monitor stand complete with transparent front lip for screens, and has a braided USB cable hardwired in. Pulling the pill-shaped camera out of the box, you'll notice how light it feels, almost like empty plastic, and the mechanical notch for the privacy shutter.

Sadly, the shutter doesn't feel satisfying to move, but it is very easy to use and feels noticeably different when it's in and out of place, so at least it's unlikely to be accidentally in the wrong position.

This camera is also completely devoid of a mic, which we think is excellent. Webcam mics are never good and are usually too far away to begin with. Elgato has thus saved a ton of frustration, be it from programs choosing it as the default mic, or having just another device listening in. No mic plus a privacy shutter means you can leave this cam plugged in and don't have to worry about yet another digital spy or accidental embarrassment. Plus the LED is large and remains on whenever the camera is active. It's all a huge step up from a tab of electrical tape over the lens.

The Elgato Facecam Neo pumps out 1080p at up to 60fps, making it more than enough for most streams and meetings. For a light little webcam that can fit easily into most settings, it's very capable. It does a great job of colors and skin tones, but tends to lean just a little toward oversaturation, which ramps up when it tries to compensate for low light.

We've plugged it into a few PCs with different software and it's been seamless every time. However, you might want to grab Elgato's free camera hub software to get access to all the features. The hub has controls for many things, like zoom and brightness, which can be saved to the camera to be used in other apps. This was also the only way we could see to enable the sought-after HDR setting.

© ELGATO

APATHETIC AUTOFOCUS

Enabling HDR will reduce video down from 60fps to just 30, but does make a huge difference if you've got sunlight messing up your shot. However, turning on HDR can make the autofocus sluggish and inaccurate. Indeed, the autofocus isn't particularly snappy in standard SDR mode. But it usually gets there within a few seconds, and then holds it well until there are any drastic changes.

There are also options to turn the focus from auto to fixed in the software, which might be a workaround if you need it. Unfortunately, all of this means you might actually need the Elgato camera hub to make changes. Being able to save settings to the camera is a huge help here, but it makes us wish it had profiles, or even just a physical button for toggling HDR. The good news is the app is free and easy to use, so this shouldn't be a huge barrier.

While wonderfully versatile, if you're wanting a dedicated webcam to make videos, something fancier with 4K support is probably a better pick. So, this is a more a general-purpose webcam that will work fairly well even if you're plagued with horrible lighting. It only takes a few tweaks to have it working nicely at a purpose lit streamer's desk as it does to have it in an awkward office. It's lightweight, small, works immediately with all kinds of software (including FaceTime on our iPad), and 1080p at 30-60fps is more than enough for streaming and video calls. For under $100, it all comes together to make the Facecam Neo a smart choice, especially if you've got harsh light or use cams on different setups regularly. –HOPE CORRIGAN

VERDICT **8**

Elgato Facecam Neo

PRETTY FACE(CAM) Super-easy to set up and use; Excellent security.

CAN'T FACE IT Feels a little cheap; Autofocus goes walkabout with HDR enabled.

$100, www.elgato.com

SPECIFICATIONS

Supported resolutions	1080p60, 1080p30 HDR, 720p60, 720p30, 540p60, 540p30
Optics	Elgato AF lens
Focus	Fixed (30-infinite)
Field of view	77°
Sensor	CMOS sensor 1/2.9-inch
Connection	USB Type-C
Dimensions	3.5x1.3x1.6 inches

Frostpunk 2 is much bigger than the first game.

CITY-BUILDING SURVIVAL GAME

Frostpunk 2

Endure a blizzard of tough choices and lasting consequences in the survival city builder sequel

CITIZENS OF NEW LONDON: As you prepare to banish us from the city forever, we hope you understand that we always had good intentions. We never meant to run out of food, let hundreds of you perish from cold, or have the streets patrolled by giant fascist robots who stomp entire buildings flat while trying to prevent a single starving child from stealing a chocolate bar. Things just got out of hand.

Thirty years after the events of the original, the world is still freezing cold and society is still a hot mess. *Frostpunk 2* is much bigger than the first game—in *Frostpunk* you finish the campaign with a city of maybe 800 citizens, while now you start with a population 10 times that. This increase in scale isn't entirely successful; as a city builder, *Frostpunk 2* is far more abstract than the original, and we never felt much of a connection to, or interest in, our city as a physical place. As a society management sim, however, *Frostpunk 2* is as effective as the original. It's packed wall-to-wall with torturous choices, agonizing consequences, and a sliding scale of morality that's as slippery as ice.

Once again, you're in charge of New London, decades after it was founded and built in the original game. The breaking news isn't great: Coal, which has kept the city's generator burning for 30 years, is dwindling. It's time to expand the borders and harvest more materials from the frozen ground outside the city, while restlessly searching the distant frigid wastelands for new resources, such as oil, to keep your people alive and warm.

In *Frostpunk 2*, you no longer place roads, houses, or single factories in your city as you expand, but lay out massive districts that can support thousands of people. Instead of a timescale of hours or days, entire weeks rush by in a matter of seconds. This transformation from the small settlements of the original to a metropolis is exciting, especially visually, but it loses something in the process.

Our city looks gorgeous, a sprawling network of dieselpunk pipelines, busy roadways, and fume-belching factories, but even at its slowest speed, it's like looking at a city in timelapse, which makes it feel all the more artificial. Even after hours of construction and planning, our city never felt like a place we built. You can zoom in on one building per district for a closer look and see a vignette of citizens at work, but it never feels like you are peering into a living, breathing place.

Influencing every action you take in *Frostpunk 2* are New London's factions and communities. You begin the game with only two or three factions in your city, but as the months and years pass, the choices you make, or don't make, give rise to more. Our city's biggest faction was the Stalwarts, law-and-order fanatics who pretty quickly started suggesting disturbing ideas like thought-control programs and human experimentation on prisoners, even before we'd built a prison.

But our choice to research and implement new tech saw the rise of a group called Pilgrims, who loathe technology so much they want us to shut off the giant generator that's kept humanity alive for the past 30 years. Hardcore survivalists called the Icebloods also appeared, extreme badasses who walk around shirtless in the cold and wrestle bears, but also think city guards should be immune from prosecution. They worry us. They all worry us.

Getting laws passed among these conflicting factions almost always requires making deals. A group can be

Factions don't just beef with you, but with each other.

Coal has kept the city's generator burning for 30 years, but now it's dwindling.

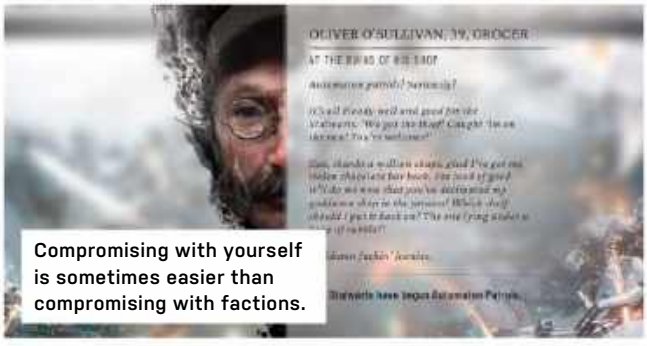

Compromising with yourself is sometimes easier than compromising with factions.

Frostbreakers smash the ice and expose the resources beneath.

persuaded to vote against their interests in exchange for some consideration later, such as agreeing to research a branch of technology they favor, or letting them set the agenda for the next vote. Break a promise you made, and their approval of you plunges, sometimes so much they become actively hostile and begin holding anti-you rallies, which can shut down your districts until you either start giving them what they want or send in a stronger faction to run them off. There's a dilemma for you: Should you squash the zealots by empowering the fascists? Spoiler: We did.

PATH OF LEAST RESISTANCE

That was just one of many steps we took down a darker path. Despite our semi-good intentions, we slowly found ourselves worn down by the endless political give and take, and the occasional major crisis, and soon we were making decisions not because they felt right but because they would give us what we wanted with the least resistance. Compromising with ourselves was sometimes easier than compromising with factions. "OK, it's gonna make the zealots even angrier if we pass an organ donation law, so we'll let them have religious funerals instead. And if people need transplants, well, we'll deal with that later. Somehow."

Once we started betraying our own values, it was pretty easy to keep on doing it. At first we scoffed at the idea of robot patrols, but then crime went through the roof and passing reasonable laws would have taken far too long and cost more political capital than we could spend. And there was our city's biggest and most favorable faction essentially elbowing us and pointing at a button labeled "Activate giant robot cops." Crime went down. That makes it worth it, right?

Thankfully, we could take a break from our massive toxic city on the brink of disaster by poking our head up and exploring the world around us. The wasteland is dotted with resource nodes to harvest, frozen settlements to plunder, and tidbits of lore about what's happened in the decades since the events of the first game. Best of all, there are larger locations in the world you can turn into proper colonies. Peel off a couple thousand citizens, send them there, and begin building housing and mining operations just like at home, though on a much smaller scale than New London.

Not so shockingly, on more than one occasion our trust meter—the constantly shrinking line at the bottom of the screen—went completely dark. Pass enough unpopular laws, break enough promises, run out of food because you were too busy tinkering with a colony instead of minding the shop, and your competing factions finally agree on something: Kicking you out of office and into the snow to die. On the plus side, there are lots of ways to start over: In addition to the New London campaign, there's a sandbox mode with seven different starting locations, different win-states to choose from, and other modifiers to play with.

As a city builder, *Frostpunk 2* is a step down from the original due to the increase in scale, which unfortunately keeps the city at arm's length. As a society sim, however, it's every bit as engrossing as the first *Frostpunk*. Like a tiny snowball rolling down the side of a mountain and eventually becoming an avalanche, even the smallest choices can have major consequences. –CHRISTOPHER LIVINGSTON

VERDICT **8**

Frostpunk 2

COLD AS ICE Excellent story mode that teaches gameplay; factions add a challenging political layer; clear communication around resource management.

WILLING TO SACRIFICE Multiple settlements feel more burdensome than challenging.

RECOMMENDED SPECS Intel i7 or AMD Ryzen 7 2.8GHz; 16GB RAM; Nvidia RTX 2060, AMD Radeon RX 5700, or Intel Arc A77.

$44.99, www.frostpunk2.com, M-rated

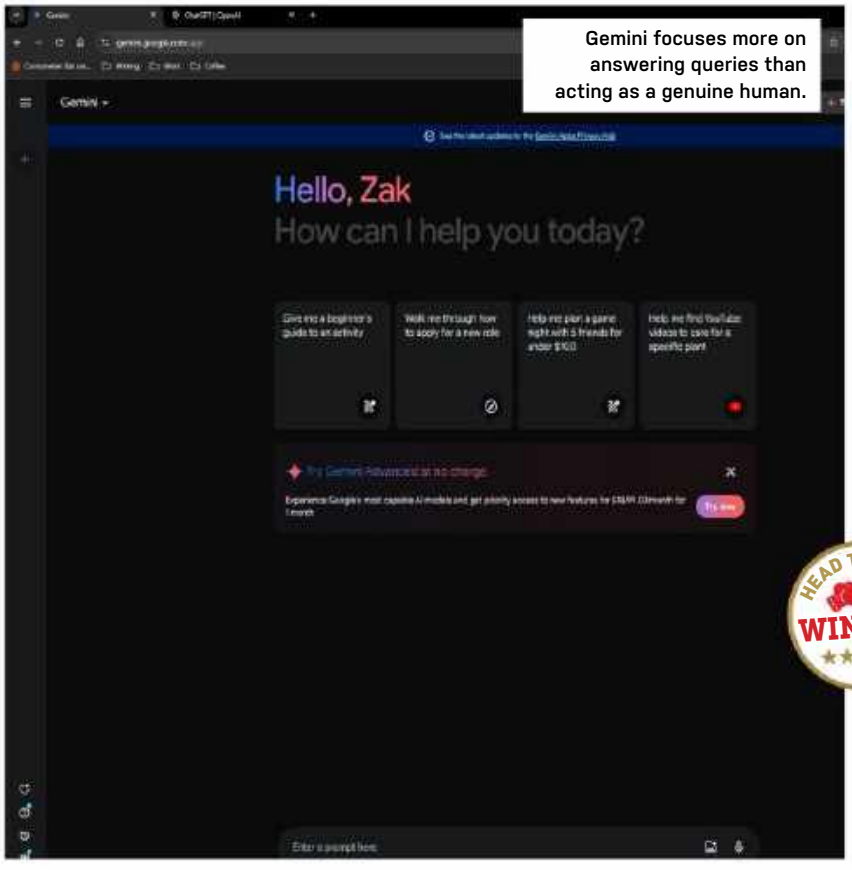

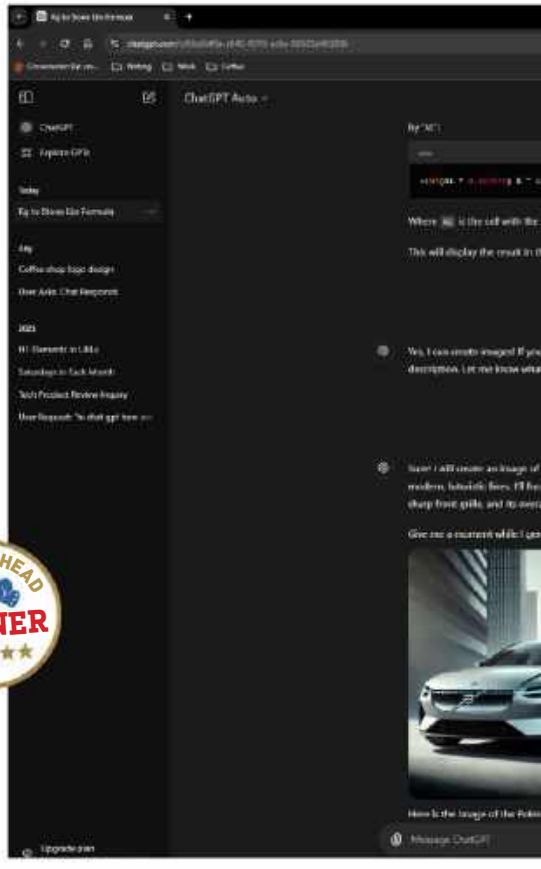

Google Gemini vs. ChatGPT

Which LLM currently holds the crown?

IT'S PROBABLY FAIR TO SAY that without the likes of ChatGPT, our current era of AI fervor wouldn't be happening. It's the language model that started the AI rush, the one that brought all manner of business concepts, products, websites, and apps to prominence in our daily lives. AI has been here much longer than that, of course, with the likes of smart speakers converting speech to text, video-gaming PvE encounters, and not forgetting DLSS and supersampling being pivotal examples of it, but it's the language models that really shook up the ethos. It's those key LLMs that brought generative AI to the forefront of the human imagination, for better or worse.

Today, we're putting two of those LLMs under the spotlight: ChatGPT and Google's Gemini. Both of them are large language models capable of an insane amount of utility, both are trained on billions upon billions of words, documents, books, and content beyond imagining, and both—at least on the surface—are free to use. So, which one has the edge? Which one out of the two large language models currently holds the LLM crown, and which is the one you should be using?

ChatGPT debuted to the world on November 30 2022, and was built by the OpenAI team. It's a large language model predominantly built around mimicking human conversation above all else. It's still capable of the vast array of content creation that generative AI is known for, but the team has prided itself on producing a chatbot that primarily responds in an incredibly similar manner to a human.

It's not a cheap thing to run, either. Microsoft initially built the team a bespoke server based on the Microsoft Azure supercomputing infrastructure. According to some sources, it currently houses well upward of 30,000 flagship Nvidia GPUs. It's an interesting yet often underestimated fact that language models actually require a far greater amount of processing power than that necessary for image generation or upscaling (as is perhaps obvious, given DLSS's localized computing success).

Gemini, on the other hand, was launched by Google back in March of 2023, initially under the name Bard. It uses an entirely different architecture than ChatGPT, in Google's Pathways. In fact, if you ask ChatGPT about it directly, it'll happily tell you that Pathways is far more capable of handling multi-modal tasks and that its access to more current data far outstrips that of ChatGPT as well. In fact, current models of ChatGPT don't have access to any information later than September 2021, whereas although not publicly confirmed, Google Gemini's knowledge cut-off is seemingly up to date as of today.

The big problem, however, is that of accuracy. It's worth remembering that both of these models have been trained in

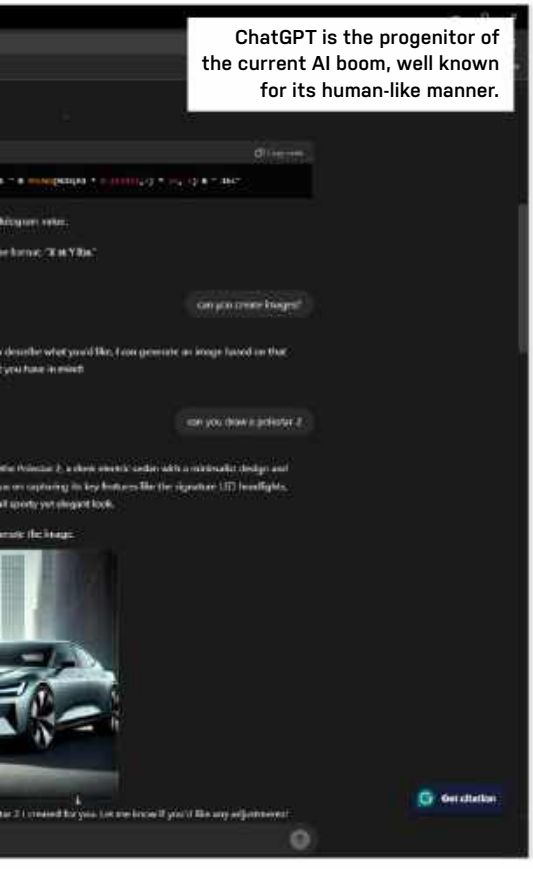

ChatGPT is the progenitor of the current AI boom, well known for its human-like manner.

© OPENAI

different manners. ChatGPT is famously (or infamously) trained via a reward model, where invigilators were hired to auto-correct all manner of ChatGPT responses, to ensure it provided the most accurate life-like responses, without falling into the realm of discrimination or causing offense.

This is very similar to rewarding a dog for sitting, and as such, does occasionally lead to responses that are less accurate than Google has managed with its own model, because ChatGPT looks to create the most human-like response, rather than the most accurate one. For instance, ChatGPT states that this journalist has been working in tech for 15 years and lives in San Francisco. That is incorrect on multiple fronts, particularly if its knowledge is limited to 2021. Ask the two bots to create an Excel formula converting metric kilograms into stones and pounds, however, and not only are both capable of producing the formula, but ChatGPT provides you with a more precise variant, including decimal places. Both are also more than capable of coding and producing Python scripts, among other problem-solving applications.

Still, Gemini is very much acting in the Google way, as a knowledge hub, rather than a more personable chatbot. One large advantage of ChatGPT is that it's now capable of storing prompts and messages as well, meaning that it's possible to have full-on conversations with it, referring to past comments in that same thread. ChatGPT is well known for having a vast number of character bots out there based on it that can serve this very purpose, although to a somewhat limited degree. As you can imagine, there have been some serious concerns raised in regard to the ethical ramifications of using these tools as virtual partners.

GENERATION & INTEGRATION

When it comes to image generation, both Gemini and ChatGPT now support this in their native browsers. Just chuck a prompt at either one, along with some keywords, and it'll quickly spit out an image for you. ChatGPT has a limit on the number of images it can create in its free edition, however, while Google Gemini currently doesn't for those on the Google One plan—and if you pay the $20 monthly fee, image generation is unlimited in its capacity. Gemini also seems to be able to produce more accurate images comparatively, but still, if you're looking for more abstract art elements, ChatGPT is definitely a sure-fire pick. We'd highly recommend using these for inspiration, however, rather than copying and pasting. Both chatbots still struggle with text in particular, and it can often seem blurry or out of place.

Integration across the two is equally becoming a larger talking point. Gemini is now more readily available across all of Google's software suite, which includes Gmail, Google Docs, and even Android devices. All you have to do once you've opted in on the Android front is talk to it as you would your old-school smart speaker, and you're immediately utilizing Gemini. ChatGPT also has access to a vast array of plugins built for specific APIs, including Zapier, Shopify, and Expedia, to name just a few, whereas Google still currently lacks that capacity. In fact, Google is entirely insular right now, with the Workspace Search and Cloud Platform being the only places you can take advantage of it, because Alphabet hasn't released the API or SDK to the public yet, unlike ChatGPT. The latter runs a pay-as-you-go system for developers, where the number of uses (or tokens) leads to incurring a cost, rather than having to pay an upfront fee.

This costing system makes a lot of sense, given the nature of LLMs and just how much processing power they require to run. It's not exactly something you can facilitate locally, even on the most powerful private server.

CONCLUSION

Both models are remarkable in what they're capable of. Of the two, Google's Gemini is incredibly accurate, both from a data standpoint and in its own generative capabilities. It's hard not to be impressed if you ask it to write a story or create an image. Yet ChatGPT has character, a more human-like nature that gives it a far more approachable feel, which is a win in itself. Additionally, that wide-reaching integration and API release has helped its uptake; although Google has the Android OS and its own software suites, it feels like ChatGPT still has the edge in regard to the number of applications it can be integrated into, which bodes well.

It's difficult to pick a direct winner. Each LLM clearly serves a different purpose. ChatGPT is designed to be a companion. It has the capacity to function like a travel manager, a coder, a creative, but at its core, it's humanity that it's trying to emulate. You can see that with how it's been trained on the data to which it has access. It can get the latest news, but it takes time and doesn't know it instantly. Google is the opposite. Gemini's built around answering queries, creating content, and generating responses on the spot. But it lacks memory, personality, and the subtle nuance that makes ChatGPT almost likable to some extent. They're both unique in that way, providing a different solution to different problems, while smoothly sliding into each other's wheelhouse as they do so. –ZAK STOREY

VERDICT 8

Google Gemini

WE ARE THE BORG Better accuracy; integrated natively into Android OS and Google workspace; image creation available on basic plan; better creativity all in all.

RESISTANCE IS FUTILE Less personable; lacks memory; best content creation behind paywall; no public API available.

$0-$20 per month, https://gemini.google.com

VERDICT 8

ChatGPT

TURING TEST? WHAT'S THAT? Phenomenally life-like; rapid responses; abstract creativity; open platform; can remember prompts; brilliant with speech.

AMPED-UP TAMAGOTCHI Accuracy still missing; imagery off the mark; paid version required for images in general; missing latest information.

$0-$20 per month, www.chatgpt.com

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> More ports needed
> Home server help

Easy USB-C

I'm wondering if you or your staff of experts can help me find a solution to a high-speed USB port shortage. I need a USB-C card that will work in a PCIe x16 slot. (Is it possible to install a PCIe x1 card in an x16 slot?) The motherboard is an Asus ROG Strix B365F Gaming. I've used the only usable x1 slot. If I understand correctly, one of the two remaining slots shares bandwidth with the second NVMe socket and the other x1 slot is physically blocked.

I want to set up an external drive that can accommodate at least 5Gb/s speed for backups. I have found an Inateck Power Supply PCIe x16 card that is 80 bucks' worth of overkill. It provides 30 watts of phone charging power and has a total of eight ports.

The computer dual boots PCLinuxOS Debian Edition and Windows 10 (the TPM has been removed to prevent downgrading to Windows 11.) The Doctor may remember this computer from the Holiday 2021 issue.

–D. Marshall

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS:

Thanks for your question—I do love a chance to go into the history of PC I/O, as well as try to help solve connectivity issues. I guess I'm just a bit of a cable nerd at heart!

First of all, it's been 10 years since USB-C was introduced to the PC market with the v1.0 specification. Most motherboards today come with at least one USB 3.2 port, which is capable of transfers at a speed of 10Gb/s, while if you opt for a pricier motherboard like the MSI MEG Z790 Godlike ($550), that features USB-C with Thunderbolt 4 offering 40Gb/s speeds. Speedy!

While premium motherboards like the one above now come with three USB-C ports on the rear, plus two more Type-C connectors on the motherboard for compatible cases, in the low-to-mid-range market, it's still much more common to only get one Type-C port on the motherboard, and another one on your case. That's the situation I have as I look down at my NZXT H7 Elite case and Z790 motherboard as I write this.

Desktop PCs don't usually have a lot of USB-C ports; this $30 adaptor can help.

It's not a major issue day to day, it just seems a little outdated when every piece of portable tech I own (laptop, mobile phone, tablet, headphones, watch) uses USB-C as standard for data transfer and charging. Anecdotally, I'd heard that it hadn't become more adopted as it's quite a flimsy little connector that can easily be damaged—I've never broken a port myself, but I admit that USB-A is sturdier.

History lesson over, I think the simplest solution to your problem is to buy a USB-C hub. Accessory manufacturer UGREEN sells a four-port USB adaptor for $30 that will give you the 5Gb/s speed you require for backups (it's a USB 3.2 hub, so actually supports 10Gb/s, but your motherboard maxes out at 3.1 speeds, which are 5Gb/s). You should note that such a device won't support charging or media display, but as you've said you want something for backups, that shouldn't be an issue. It's also compatible with both Windows and Linux.

If you're dead set on going down the PCIe card route, then you can, although as you note, it will be more expensive. To answer your question, it's no problem at all to put a PCIe x1 card in an x16 slot. StarTech.com does a two-port USB 3.2 card for $47, which will offer you full 10Gb/s transfers, the downside being it only comes with two USB-C ports. But again, that should be enough, and you

↘ submit your questions to: editor@maximumpc.com

might be able to justify the extra $17 over the UGREEN option given it offers double the speed.

Finally, I note that you're using a five-year-old motherboard using the LGA1151 socket for 8th and 9th-gen Intel Core processors. I'm sure it's still performing well and that you don't want to upgrade your system, but it is, of course, another option. You'd need a high-end motherboard like the MSI one listed above to get more than one USB-C port, but on the plus side, they go up to four times faster than the add-on card listed above. These higher-end mobos will also support cases with more than one USB-C port on the front panel, such as the Asus ROG Hyperion GR701 (featured on our September 2024 cover), but again, a case such as this carries a hefty premium at $500. Probably not the option you'll take given the $1,000-plus cost, but I like to be thorough!

Server solutions

I'm not sure if you've done this lately, but I'm thinking about what it would take to build a media and file server. I do some electronic music production; drum and bass is my genre. I have a few synths and modular gear that I record from constantly and I end up with terabytes of WAV files.

I also have quite a bit of other media stowed away. My current solution is a TrueNAS server with about 16 terabytes of spinning HDDs in it. I've been moving a lot of files recently and this has been painfully slow. That got me thinking, maybe it's time for another server upgrade.

One thing that I would like to think about is using SSDs or even NVMe drives. I know it's possible to get PCIe cards that they slot into. I'm thinking I would want 20 or 30 terabytes minimum. I also need to consider the speed of the network access to it, including what Ethernet cables I need to use and what switch I need to complete a fast connection. I'm willing to spend a couple of thousand dollars on this to get a better solution in place.

One last thing: It would be nice to have a way for people outside of my network to be able to access some files. People that I collaborate with regularly. Currently, I try to share files through Dropbox, but if there is a solution to allow sharing through the NAS, that would be awesome. I've learned to use TrueNAS well enough to set up pools and 'partitions' with Windows SMB sharing. I'm not opposed to a better solution, if there is one. If you have a recent article that I've missed, feel free to point me toward it. Thanks for any advice.

–S. Rakestraw

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS:
It's always good to hear from a fellow NAS lover! Personally, I've owned a few different types, starting with an HP Microserver running Windows Home Server in around 2013, to more recent off-the-shelf solutions from Synology and QNAP. Four years ago, electricity costs had increased substantially in my area, so I gave my NAS up in favor of moving completely to the cloud, but I just missed the flexibility of running my own server too much and got back into the NAS life again.

Personally, I run a Synology solution (the Synology DS423+), which I would highly advocate for. I know, it's not very *'Maximum PC'* and I always read build-your-own NAS guides with interest and dream of giving it a go, but I personally find networking can be a bit of a timesink, and I appreciate how easy ready-made solutions make it. I've also needed technical support on more than one occasion, and considering this is about my precious data, I do think it's worth paying a bit more to have that technical expertise to hand if it's needed. They also make data redundancy super-easy to set up, are low energy out of the box, run (fairly) quietly, and are very compact.

There are downsides, however, the main one being that my particular model is still only offering 1GbE Ethernet, although there are cheap 2.5GbE USB adaptors available. It's also expensive to expand beyond four drives, and it's pretty underpowered in terms of what I could build myself for the same amount of money.

If you want to give it a go and upgrade your existing server, then your first port of call should be Nick Peers' feature "Build Your Own NAS 2024" in the April issue of the magazine. In it, he puts together a small and power-efficient build that uses a fast SSD as a boot drive for just $565—a bargain, really. As Nick says, a build like this is far more flexible and powerful than a pre-built option, and should therefore last a lot longer as a result.

In terms of storage, I still personally stick to mechanical hard drives just because of their unbeatable value, but most servers these days do have the option of using SSDs if you prefer. Or you could use one to store the files you need to be able to access quickly—the WAV files you mention, for example.

In terms of cabling, make sure you pick up CAT6 Ethernet cables if you have devices that go above 1Gb/s—CAT6 goes all the way up to 10Gb/s. As you mention, you'll also need a faster network switch or router if you're keeping everything wired up—these tend to be a lot more expensive than 1Gb/s options, which is why I've personally never upgraded.

Finally, in terms of sharing files, it's pretty easy to create links from TrueNAS or Synology. My ISP doesn't offer a static IP sadly, but I found Tailscale, which is essentially a VPN that gets round dynamic IP, to work perfectly. I also always have my NAS backing up important files like documents and photos to online storage such as OneDrive and iDrive.

Guy recommends the Synology DS423+, his personal NAS solution.

THE BUILDS

THIS MONTH'S STREET PRICES...

WE'RE FINALLY INTRODUCING new processors into our Blueprints mix. A lot has happened over the last month. We've had Intel dealing with massive stability issues, both at a physical and a software level, along with AMD launching new flagship chips riddled with latency and Windows bugs. But it finally feels like things are calming down, and everything's going back to normal. The latest microcode updates are now available for both AMD and Intel, with AMD's significantly reducing latency on its 9000 series, and Intel massively curtailing power draw, without sacrificing too much on the performance front. With that, we've decided to include new chips across a swathe of our builds.

First up is the budget end of the spectrum. Both systems below clock in at around the $900 mark. We've gone back to the Core i5-14400F for Intel, which has bumped the CPU price up by $4, although it's worth pointing out that our last-gen pick featured a CPU cooler bundled in with that price. Still, the 14400F does pack in 10 cores, six of them performance and the other four efficient cores as standard, along with 16 threads, and a comfortable 65W TDP to back it all up. Plus it comes with a fairly basic stock cooler.

We've also had to change out the memory, opting to go for the same kit as in our AMD build. And, sadly, our SSD from last month is no longer in stock, but fortunately WD's SN770 500GB PCIe 4.0 drive is, and at a super-low $45, perfect for an OS.

Our AMD system generally didn't see any major changes outside of that SSD swap. Pricing remained fairly stable across the board, In fact, we actually saw a price drop on the MSI PRO B650-S mobo there, down by $20 from last month. Memory pricing also continues to fall, and we're seeing even more 32GB kits at C32 6400 land with us now. Both CAS latency and frequency are improving, along with prices falling, which is great news for us as consumers.

AMD INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$90
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$62
Mobo	MSI PRO B650-S Wi-Fi	$130
CPU	AMD Ryzen 5 7600	$182
GPU	ASRock Challenger D Arc A750 8GB	$190
RAM	32GB (2x 16GB) Silicon Power XPOWER Pulse Gaming @ 6400 C32	$87
SSD 1	500GB Western Digital Black SN770 PCIe 4.0 M.2 NEW	$45
SSD 2	1TB Kingston NV2 PCIe 4.0 M.2	$57
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $875

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 i5-14400F NEW	$200
GPU	ASRock Challenger OC RX 7600 8GB	$250
RAM	32GB (2x 16GB) Silicon Power Value Gaming @ 6000 C30	$87
SSD 1	500GB Western Digital Black SN770 PCIe 4.0 M.2 NEW	$45
SSD 2	1TB Kingston NV2 PCIe 4.0 M.2	$57
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $921

BOTH OF OUR MID-RANGE systems have taken a hit on the pricing front. Our AMD build has gone up by $133, and swapping back to the Intel Core i7-14700F increased our blue rig's total cost by $91. The chips were the major drivers of this. Initially, in our AMD system, the board we had last issue, the Asus Prime X670-P Wi-Fi, had bounced up to $244. We swapped that out for the ASRock X670E Pro RS to save us $33.

It's worth noting that, depending on the motherboard BIOS, you may have to update the BIOS to the latest AGESA code. AMD's currently running AGESA 1202; not only does it add support for the Ryzen 9000 series (as does the one prior), but it also reduces inter-core latency by around 60 percent, which is a huge win as gaming performance has been on the sluggish side since launch.

We've also swapped out the memory again, going for a peppier Teamgroup T-Create Expert 32GB kit at 6400 C32. Interestingly, this is still only $95, and the 64GB kits are now starting to enter the $150 area, so if capacity is your thing, it may be worth spending a little extra there. Otherwise, we saw some fairly steady price increases across the board, both SSDs saw a $1 or $2 increase, the 7700 XT likewise had a $5 bump, and all this has led to an overall increase of around $133 on last month.

For our Intel system, changes have similarly been minimal outside of that processor shift. We've had to shuffle around the GPU, as last issue's Zotac edition had climbed up by $40 or so. Fortunately, MSI's Ventus is actually one of the cheapest RTX 4070s we've ever seen at $520. Similarly to our AMD system, we recommend updating your motherboard's BIOS, regardless of whether the mobo supports the chip, just to ensure that Intel's latest microcode update is working, and to limit any chance of system instability occurring later down the line. If that happens, it results in permanent damage to your processor. Not good.

AMD INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow 2022	$100
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold	$95
Mobo	ASRock X670E Pro RS NEW	$210
CPU	AMD Ryzen 7 9700X NEW	$337
Cooler	Noctua U12S chromax.black	$90
GPU	Asus Dual OC Radeon RX 7700 XT 12GB	$400
RAM	32GB (2x 16GB) Teamgroup T-Create Expert @ 6400 C32 NEW	$95
SSD 1	1TB Lexar NM790 w/Heatsink M.2 PCIe 4.0	$82
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:		**$1,545**

INTEL INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow 2022	$100
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold	$95
Mobo	MSI Z790-S Wi-Fi ATX	$150
CPU	Intel Core i7-14700F NEW	$347
Cooler	Enermax Liqmax III 360 ARGB 360mm AIO	$80
GPU	MSI Ventus 2X OC RTX 4070 12GB NEW	$520
RAM	32GB (2x 16GB) Silicon Power Xpower Zenith Gaming @ 6000 C30	$79
SSD 1	1TB Lexar NM790 w/Heatsink M.2 PCIe 4.0	$82
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:		**$1,589**

WE KNOW what you're thinking: Ryzen 7 9700X in Mid-Range, Ryzen 9 9950X in Turbo, right? Not today. That's not how we do it. Pricing on the X870E boards is up in the air, selection is slim as they're still making their way into the market, but until then, we're going to hold off on committing to that 9950X. Don't get us wrong, it's a potent processor, and its multi-core performance is off the chain, but to really take advantage of it, it needs to be in an X870 motherboard. Additionally, right now, the 9950X is clocking in at around $560, which is good value (particularly given Intel's 14th gen flagships are still at that price point), but because the 7950X is $450 and still provides some top-tier performance, it's the smarter pick at the moment.

For that reason, our Turbo build for AMD remains relatively stable, with the only major change coming in the form of the MSI MAG X670E Tomahawk Wi-Fi motherboard. It's available for just $260 and is $20 cheaper than our board from last issue. We've also seen price drops, albeit quite small ones, with the 7900 XTX dropping by $10, and the 64GB Patriot Viper Venom falling by $5. Sadly, however, the Enthoo Pro 2 has seen a price increase of $11. In total, our AMD system has gone up by $8, which is a shame, but not the end of the world, and likely a lot less than it will do next issue with the arrival of the X870 boards.

On a similar note, as we don't have the mad cash savings on our Intel system anymore, thanks to 12th gen, moving back to the 14900KF (the cheapest 14900 we could find) has increased the overall CPU price by $227. Ouch. To compensate slightly, we've gone for the Asus Prime Z790-A Wi-Fi motherboard instead of the Gigabyte Z790 Pro X ATX we had last month. That does save us a bit of cash, albeit just $30, but it does also deprive us of a bit of connectivity as a result. Otherwise, similar to the AMD build, pricing has remained relatively stable again, although we did see a $1 increase on the AIO. Interestingly the PNY Verto RTX 4080 Super at $960 still remains the best value 4080 you can get, which is some good news. It is a shame as the RTX 4080 Super was meant to be a $999 card, but, well, manufacturers....

AMD INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$156
PSU	Super Flower Leadex VI Platinum Pro 1000W - 80+ Platinum	$130
Mobo	MSI MAG X670E Tomahawk Wi-Fi AM5 **NEW**	$260
CPU	AMD Ryzen 9 7950X	$449
Cooler	NZXT Kraken 360—360mm AIO	$180
GPU	ASRock Phantom Gaming OC Radeon RX 7900 XTX 24GB	$900
RAM	64GB (2x 32GB) Patriot Viper Venom @ 6400 C32	$175
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:		**$2,617**

INTEL INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$156
PSU	Super Flower Leadex VI Platinum Pro 1000W 80+ Platinum	$130
Mobo	Asus Prime Z790-A Wi-Fi **NEW**	$250
CPU	Intel Core i9-14900KF **NEW**	$507
Cooler	Asus ROG Strix LC II ARGB360mm AIO	$134
GPU	PNY Verto Overclocked RTX 4080 Super 16GB	$960
RAM	64GB (2x 32GB) Patriot Viper Venom @ 6400 C32	$175
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:		**$2,679**

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US LLC, 130 West 42nd Street, 7th Floor, New York, NY 10036. USA. Website: www.futureus.com. Future US LLC also publishes *MacLife*, and *PC Gamer*.

UNEARTH THE FULL STORY BEHIND THE LEGEND OF ZELDA

Explore every game in the series and what makes each one so great. If you're a fan, you'll enjoy the trip through time. If you're a newcomer, we offer you this sage advice: "It's dangerous to go alone, take this!"

ON SALE NOW

FUTURE

Shop easily online at:

magazinesdirect.com

bit.ly/3xbOGXq

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside