

MSI MEG Z690 UNIFY
First taste of Intel's
next-gen platform **PG. 80**



GAMING ON THE GO
XPG's awesome Intel-
designed laptop tested **PG. 76**



ULTIMATE GEEK QUIZ
Know your storage protocols
from your CPU SKUs **PG. 26**



MAXIMUM PC

MINIMUM BS • JANUARY 2022 • www.maximumpc.com

2022 TECH PREVIEW



**BEST GAMING
CPU EVER!**

Intel's awesome
Core i5-12600K
reviewed **PG. 68**

THE FUTURE OF THE PC REVEALED:

- ✓ Killer next-gen CPUs
- ✓ 3D deathmatch: AMD vs. Nvidia vs. Intel
- ✓ Mini-LED monitors & more



**ROUTER TO
PERFECTION**

Everything you need
for a super speedy
home network **PG. 34**



**LINUX
LOVE FEST**

How Windows and
Linux have become
best buddies **PG. 34**

DUNE



"AN AWE-INSPIRING MASTERPIECE."

—FOX TV

LOOK FOR IT ON
4K ULTRA-HD™, BLU-RAY™ AND DIGITAL

SCAN TO SEE MORE



BECOME AN EXPERT CODER THE EASY WAY



FUTURE



Ordering is easy. Go online at:

magazinesdirect.com



Or get it from selected supermarkets & newsagents

inside

JANUARY 2022

SUBSCRIBE TODAY!
PG. 32



QUICKSTART

8 THE NEWS
AMD vs. Intel heats up; Facebook stops spying; Nvidia upscaling tech

11 TECH TALK
Could multi-chiplet GPUs be the future of computer graphics?

12 THE LIST
The top six VR headsets



Featuring the best VR headsets from Oculus, PlayStation, HTC, HP, and Valve



© MONSIEUR/JETTY IMAGES, GOOGLE, BRAVE

R&D

51 HOW TO
Draw patterns in Illustrator; build a better file server on Raspberry Pi; get Steam gaming on Linux

52 AUTOPSY
We tear down Apple's new MacBook Pro 16-inch to find out how easy it is to repair

LETTERS

14 DOCTOR

94 COMMENTS

26 ANNUAL GEEK QUIZ

Have you been paying attention? Time to find out with our annual geek quiz

34 LINUX AND WINDOWS 11

Windows and Linux living together in perfect harmony... or is this the perfect storm?

42 SECURE YOUR ROUTER

Our guide to adding extra security layers to protect you and your network

IN THE LAB

70 ASROCK RADEON RX 6900 XT FORMULA



80 MSI MEG Z690 UNIFY

83 ASUS ROG STRIX GO CORE MOONLIGHT WHITE



92 GOOGLE CHROME VS. BRAVE



89 AVERMEDIA PW513





PAPER REVOLUTION

Over 66% of paper and paper-based packaging in the U.S. is recycled into new products, one of the highest recycling rates of any material in the country.

Discover the story of paper

www.lovepaperna.org

Sources: American Forest and Paper Association, 2019
U.S. Environmental Protection Agency, 2017

Love Paper is a registered trademark for Two Sides Ltd. Registered in the U.S., UK and other countries and used with permission.



Join the conversation!



Discuss all things PC gaming with fellow readers at our official forums
forums.pcgamer.com

PC GAMER

MAXIMUM PC

EDITORIAL

Editor: Jeremy Laird

Staff Writers: Christian Guyton, Sam Lewis
Contributing Writers: Alex Cox, Ian Evenden,
 Chris Lloyd, Nick Peers, Jarrod Walton
Production Editor: Jon Crampin
Editor Emeritus: Andrew Sanchez

ART

Art Editor: Fraser McDermott
Photography: Phil Barker, Oly Curtis, Neil Godwin
Cover Photo Credits: Getty Images, Intel

BUSINESS

US Marketing & Strategic Partnerships: Stacy Gaines,
 stacy.gaines@futurenet.com

US Chief Revenue Officers: Mike Peralta,
 mike.peralta@futurenet.com

East Coast Account Director: Brandie Rushing,
 brandie.rushing@futurenet.com

East Coast Account Director: Michael Plump,
 michael.plump@futurenet.com

East Coast Account Director: Victoria Sanders,
 victoria.sanders@futurenet.com

East Coast Account Director: Melissa Planty,
 melissa.planty@futurenet.com

East Coast Account Director: Elizabeth Fleischman,
 elizabeth.fleischman@futurenet.com

West Coast Account Director: Austin Park,
 austin.park@futurenet.com

West Coast Account Director: Jack McAuliffe,
 jack.mcauliffe@futurenet.com

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

MANAGEMENT

MD Tech Specialist: Keith Walker
Group Editor-in-Chief: Graham Bartow
Group Art Director: Jo Gulliver

PRODUCTION

Head of Production: Mark Constance
Production Manager: Vivienne Calvert
Project Manager: Clara Sanchez
Production Assistant: Emily Wood

FUTURE US, INC.

11 West 42nd Street, 15th Floor, New York, NY 10036, USA
 www.futureus.com

INTERNATIONAL LICENSING & SYNDICATION

Maximum PCs 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 February 1st, 2022

© 2022 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.

Future US is a public company listed on the NASDAQ stock exchange (NYSE: FUTR).
 Chief Executive Officer: www.futureus.com
 Chief Financial Officer: www.futureus.com
 Chief Marketing Officer: www.futureus.com
 Chief Operating Officer: www.futureus.com
 Chief Technology Officer: www.futureus.com
 Chief Legal Officer: www.futureus.com
 Chief Human Resources Officer: www.futureus.com
 Chief Sustainability Officer: www.futureus.com
 Chief Diversity Officer: www.futureus.com
 Chief Compliance Officer: www.futureus.com
 Chief Information Security Officer: www.futureus.com
 Chief Privacy Officer: www.futureus.com
 Chief Risk Officer: www.futureus.com
 Chief Sustainability Officer: www.futureus.com
 Chief Diversity Officer: www.futureus.com
 Chief Compliance Officer: www.futureus.com
 Chief Information Security Officer: www.futureus.com
 Chief Privacy Officer: www.futureus.com
 Chief Risk Officer: www.futureus.com
 Chief Sustainability Officer: www.futureus.com



Jeremy Laird

RETURN OF THE KING

INTEL IS BACK. Of that, there can be no question. This month, we get our paws on the 12600K, the gamer's choice among Intel's new 12th-generation Alder Lake chips. If we had doubts last month about the range-topping Core i9-12900K SKU, thanks to some scary power consumption figures, that's largely mitigated by the new Core i5 model. The 12600K reduces 12900K's 240W power guzzling problem down to a much more manageable 150W or less.

More to the point, the 12600K is now the weapon of choice for gamers. Job done for Intel? For now, perhaps, but Intel still has plenty to prove. Alder Lake is built on the so-called Intel 7 node, otherwise known as 10nm and still years behind schedule. As you'll discover in our Tech Preview feature, AMD is tooling up for its next-gen Zen 4 processors based on TSMC's 5nm process.

I suspect Intel will struggle to immediately respond. After all, that 10nm node is only now rolling out on the desktop despite originally being slated for 2015. It's easy to forget the first 10nm Cannon Lake mobile CPUs appeared in small numbers back in 2017. As for Intel's next process, the node formerly known as 7nm and now branded Intel 4 was supposed to go live in 2017. Last summer it was delayed again. We're not expecting to see the first Meteor Lake 7nm desktop chips until 2023.

When you consider the huge size of the Golden Cove cores in the new 10nm 12600K and 12900K Alder Lake CPUs, the suspicion is that Intel may have intended them for 7nm, just as Intel admits the Cypress Cove cores in last-gen Rocket Lake CPUs were "backported" from 10nm to 14nm. If so, it explains how Intel has managed to create a supposedly cutting-edge CPU that consumes quite so much power.

Whatever happens with CPUs, there's plenty to look forward to, including PC monitors with clever miniLED backlights, super-fast PCIe 5.0 SSDs and so much more. Dare we dream of graphics card prices returning to normal over the next year? There are so many factors in play, not least the reality that the world of chip production generally doesn't look like recovering from the current extraordinary market conditions—driven by the pandemic and distorted by crypto-currency mining—any time soon.

However, one thing that could help involves our old friend Intel and the competitiveness of its Arc graphics. Intel doesn't need to produce a world-beater. But if the Arc family of GPUs is competitive with the likes of, say, an Nvidia GeForce RTX 3070 or AMD Radeon RX 6800, there will be three players offering proper gaming graphics. Nothing brings prices down like healthy competition. Here's hoping.

I also hope you enjoy this issue. It's a great pleasure to be back at the helm of *Maximum PC* for a few months while Zak's permanent replacement is identified. Back in 2015, when I last edited this excellent magazine, the go-to narrative among many industry observers was the PC's inexorable slide towards irrelevance and, ultimately, oblivion. So much for that! It's a genuine joy to find both the PC and this magazine in rude health.

Jeremy is Maximum PC's acting editor-in-chief. He's been writing about PCs for 20 years and enjoys nothing more than a serious dissertation on the finer points of input lag or overclocking.

submit your questions to: editor@maximumpc.com

THE NEWS

The empire strikes back

But AMD's rebel alliance is ready for Intel's Alder Lake

INTEL HAS released the best of its new Alder Lake chips first, the headline-grabbing performance models. But what about the more affordable CPUs, the chips we stand a better chance of buying? Thanks to a slip up at Dell, among others, details of further Alder Lake SKUs went online well before the official launch. As expected, we have Core i7-12700, i5-12600, 12500,

There are two versions; Alder Lake-P with a PBP (or whatever metric Intel uses) of up to 45W and carries up to six performance cores, and eight efficient cores, along with DDR5, and PCIe 5.0 support. Alder Lake-M is the low-power version, which needs just 9W. Intel isn't resting there, with talk of the 13th-generation chips, code name Raptor Lake, arriving next fall. It will be

After this comes the main event, the 14th-generation Meteor Lake. Intel already has prototypes of this chiplet design running. This uses the Intel 4 process and is constructed using chiplets, or tiles as Intel calls them, with a processor tile (Ocean Cove), a GPU tile, and an I/O tile.

The test chip also has a mysterious fourth tile (AI anyone?). The tiles use second-generation Foveros technology to stack other chips on top, making the complete package a 3D hybrid tile design. There are no solid numbers, this is still early days. It's expected to arrive in late 2022 at the earliest and will be a big step.

Meanwhile, over in the red corner, AMD's response is a new chip architecture, and Zen 4 will also be ready next year. Its CEO, Lisa Su, gave us a glimpse of the Zen 4 roadmap at its 'Accelerated Data Center' event, which was all about servers. We learned about the 96-core Genoa and the 128-core Bergamo Epyc chips. Zen 4 will use a 5nm process which offers twice the density of 7nm, with a roughly 25 percent performance hike, so the exact IPC gain will certainly be healthy. Expect the desktop Zen 4 (Raphael) before the end of 2023, some leaked documents indicate it will have twice the L2 of Zen 3.

More immediately, we have V-Cache coming to a Zen 3 chip, probably after the holidays. It's a practical, and wonderfully simple solution: just drop a fat slice of L3 on top. AMD claims it bumps game performance by about 15 percent, which is equivalent to an architecture change if that is carried across all applications.

The patents AMD has taken indicate that V-Cache is heading to GPUs too, they are mentioned twice as often as CPUs in the paperwork. Then we have a 'fork' of the Zen design, Zen4D, where 'D' is for dense, a stripped-down version with a redesigned cache layout aiming to reduce the die size. After that comes Zen 5, AMD's first hybrid design with new Zen 5 performance cores, and employing Zen4D as the efficiency cores. There are rumors of an IPC bump of around 30 percent. Yikes.

For a while, it looked as if only AMD was making real strides. We moved from the original Zen through to Zen 3 in just 44 months. Intel appeared to be dragging its heels for a few years but now it too is sprinting. Both companies will have powerful, innovative 3D hybrid chiplet designs by 2023. Zen 5 against Meteor Lake is going to be quite a tussle. We wanted a race, and it looks like we've got one. **-CL**



AMD's CEO, Lisa Su, reveals part of the Zen 4 roadmap as the company moves to a new core and a new process next year.

and 12400, and at the bottom the Core i3-12100. Curiously, the low-end chips don't appear to have any efficiency cores. An early benchmark of the i5-12400, with six performance cores using Adobe Premiere Pro, shows it stands toe-to-toe with a Ryzen 9 5600X. Alder Lake is proving to be a match for Zen 3 across the range.

What's next? Alder Lake mobile comes first, which has entered production, ready for launch around the holidays.

socket-compatible with Alder Lake and features a 'new' hybrid design.

Parts of Alder Lake aren't ideal, the die size of the Golden Cove cores for start. The Gracemont efficiency cores aren't petite either, about half the size of a Zen 3 core. Raptor Lake is a development of Alder Lake with a tidier, streamlined design and a possible small increase in the core count. We expect a typical IPC jump of about 10 percent.

FACEBOOK QUILTS LOOKING AT FACES

YOUR FACEPRINT WILL BE DELETED

META, FACEBOOK'S PARENT company, is to end one of its most contentious practices: facial recognition. It will also delete the one billion 'faceprints' it collected from the one-third of Facebook users who opted in. As a useful dataset, it must have been valuable, but the move represents "one of the largest shifts in facial recognition usage in the technology's history," according to the VP of the AI division. It's not clear what prompted the change of heart, the company has been under fire over its attitudes to privacy, and admitted that there are 'many concerns' about the unrestricted use of facial recognition. A couple of useful features will go along with it, the audio description that names people in pictures (handy for the visually impaired), and you'll no longer be notified if your face appears in other people's media. **-CL**



GRAPHICS DUOPOLY ALL OVER

Rivals eye the lucrative GPU market

REMEMBER MATROX? How about S3 or even 3dfx? Back in the 1990s, there were dozens of graphics card designers and builders. But since then, they have either run themselves into the ground, retreated into niche markets, or been bought out. For years, there has effectively been a duopoly; Nvidia and AMD control roughly 80 and 20 percent of the market respectively.

Recent times have been harsh for the industry, losing cards to cryptocurrency miners and a chip shortage that squeezed supplies and raised prices. However, the growth of cloud services, AI, and other intensive HPC applications has fueled a resurgence in GPU development. In 2019, there were 11 developers, this year it has grown to 19. Most of this development won't find its way out of the server racks, but the most obvious new contender is Intel, which started developing its discrete graphics card business in 2017. Early next year, we will see the first fruits: the Arc Alchemist (the DG2 as was). There will be four variants, the beefiest being the DG-512EU—sporting 512 EUs, this is aimed squarely at GeForce RTX and Radeon RX territory. Intel has hinted at a price comparable with a GeForce RTX 3080, so fingers crossed for decent numbers.

Meanwhile in China, Xindong has previewed its Fenghua No.1 graphics card, and the spec sheet reads well: GDDR6X memory, support for the latest APIs, including Vulkan. The company is aiming the GPU at everything from data centers to desktops. Another Chinese company, Zhaoxin, has demonstrated a prototype discrete card, only a 75W one so far, but it's still early days.

Nvidia and AMD finally have competition and it can't come soon enough. We need more cards now and anything that help push down prices is good news. **-CL**



Could the AMD and Nvidia duopoly be over?

MICROSOFT PUSHES EDGE

New Win11 build getting obstructive

A PREVIEW BUILD of Win11 has pulled a sneaky trick by stopping a popular 'helper' app called EdgeDeflector from working. This could detect the hard-coded links to Edge that Win11 likes to scatter across the OS, then redirect them to your preferred browser. The planned update kills it dead, ensuring that Microsoft's browser pops up instead. Trying to set a rival as your default isn't easy either. When you install a new browser and open a link, you have to select "always use this app" to make your choice the default. But you only get one chance, if you miss it you'll have to set it as the default for all files or link formats individually. Only Firefox manages to successfully navigate the pitfalls to install itself as the default. You can expect Microsoft to push Edge, but this is annoying at best. We've been here before, of course, back in the days of Internet Explorer. It eventually ended in court to disconnect the OS from the browser and it seems Microsoft is going back to its old ways. **-CL**

Tech Triumphs and Tragedies

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

TRIUMPHS

GAMING FREEDOM

Windows allows you to choose where to install games from Windows Store or Xbox Game Pass: no more locked folders.

THROWN INTO SPACE

SpinLaunch plans to 'throw' satellites into space using a huge centrifuge operating at high-speed in a vacuum. A small test version works, at least.

THE ORIGINAL APPLE

An Apple-1 sold at auction for \$500K complete with cassette adapter, and koa wood case: one of 200 hand-made machines.

TRAGEDIES

STOLEN RTXs

A truckload of GeForce RTX 30-series cards, including 3090s, was hijacked en route to a warehouse in California.

RECORD WITHDRAWN

A world-beating 8GHz Alder Lake overclock by Gigabyte was rejected due to a bug in CPU-Z that can give the wrong clock speed with Alder Lake chips.

LOCKED FRAMES

Nvidia's cloud gaming service locks frame rates of 12 games to below 60fps, causing jitters, even for Priority Members.



HARD DISKS TO GO PCIE

Seagate begins move to faster data transfer

IF YOU REALLY WANT to shift great gobs of data inside a PC you use PCIe. The latest generation of processors support version 4.0, which can shift nearly 2GB/s on a single lane. There's almost nothing that you can plug into it that will hit the buffers. When NVMe appeared, it was the obvious place to go for performance, it makes SATA III look positively pedestrian at 600MB/s.

Seagate has taken the logical step and demonstrated the first commercial hard drive that can connect to PCIe using the NVMe protocol, which added HDD support this summer. The result is an experimental drive with SAS, SATA, and NVMe controllers that works 'seamlessly', according to Seagate. Much of the impetus for the move comes from data centers, where NVMe is increasingly king, and moving everything to a single interface would be simpler, and cheaper.

Do spinning discs need a new bus on a home rig? Not yet, even the fastest still can't saturate SATA III, but it is getting closer by the year. The multi-actuator drives we've been promised for a while, and other improvements will outgrow their current home. Switching to a ready-made interface with acres of headroom is an obvious step. It's at the experimental stage for now, but 'key customers' will be getting the new drives next fall. It is expected to trickle down from the servers to the desktop as prices reach parity with the existing kit. SATA isn't going anywhere just yet—it's useful if you want to connect a whole stack of drives, but its days as the primary interface are looking numbered in the long-term. Witness the dearth of news on the possibility of SATA IV. —**CL**

Nvidia open-sources upscaling tech

NO NEW CARDS, BUT we've had a few software updates from Nvidia. Most interesting is that the new version of Nvidia Image Scaling (NIS) is going open source. Nvidia Image Scaling does the same job as DLSS without the AI, which means you don't need a fancy latest-gen Nvidia card. This version uses a six-tap filter with four-directional scaling and adaptive sharpening, it doesn't include any anti-aliasing, that's down to the game. As with DLSS, NIS squeezes extra frames out of higher resolutions, so you can set a game to 1440p, and let NIS crank it up to match your 4K monitor, and it should get close to the same frame rates as running in 1440p flat. Going open-source effectively means letting it run on AMD and Intel graphics. That sounds great, but without direct game support (none have this) it will require driver support. Until then, applications such as Lossless Scaling have already added NIS support.

DLSS has also received an update to version 2.3. Nvidia says it makes "smarter use of motion vectors to improve object detail in motion, particle reconstruction, ghosting, and temporal stability". The new version has already been added to a slew of AAA titles, including *Crysis 2/3*, *Rise of the Tomb Raider*, *Cyberpunk 2077*, and *Farming Simulator 22*. Nvidia has also released its Image Comparison and Analysis Tool (ICAT) as a free download. It puts up to four screenshots or videos side-by-side for direct comparison, making it easier to see what's going on as you play with the settings. —**CL**



A neat side-by-side of DLSS, Nvidia has released a utility so you can create these at home. More importantly, it has open-sourced its Nvidia Image Scaling engine.

Haswell kills DX12 games

Still running one of Intel's Haswell chips and playing DX12 games using integrated graphics? Then it's game over—the latest driver has pulled DX12 support. For anyone still needing support, Intel suggests reverting to an older driver. Haswell was a 4th-gen chip introduced in 2013, the range included a couple of tasty eight-core 'enthusiast' chips and clock speeds of up to 4.4GHz. It lived on after Skylake, and you could still buy a Haswell-based Pentium Gold earlier this year. We have become used to the rapid pace of technology, but eight years from cutting-edge to obsolescence is still impressive. —**CL**

Qualcomm backs ARM

Last year, Qualcomm bought Nuvia, a company founded by former Apple chip designers, for \$1.4bn. With its own ARM-based Snapdragon chips making little headway outside smartphones, Qualcomm quickly got to work on new ARM designs. Apple has shown that a well-executed ARM design can work and Qualcomm has stated that a move to ARM-based chips is "inevitable". That may come as news to Intel and AMD, who are squeezing every drop out of x86 designs. The aim is to have samples of the next-gen Snapdragon chips ready by the summer, with commercial sales of notebooks following in 2023. —**CL**



Jarred Walton

TECH TALK

Multi-Chiplet GPUs could be the future of graphics


USING MULTIPLE GPUS—or any processor—to accelerate a task by breaking it up into smaller, independent pieces isn't exactly a new idea. 3dfx did that with its Scan Line Interleave (SLI) technology back in the late 1990s, and Nvidia's SLI (Scalable Link Interface) and AMD's CrossFire have been doing it since 2004 and 2005, respectively.

To date, all implementations of multi-GPU rendering have had drawbacks that prevented ideal scaling. The biggest hurdle is figuring out how best to break up the task. GPUs already split workloads up into thousands of queues of instructions running in parallel, but splitting that across multiple GPUs separated by a relatively slow interface limits scaling. Consider the RTX 3090, it has an internal memory bandwidth of 936GB/s. The current NVLink connector only provides up to 112GB/s—about one-eighth the bandwidth.

Games generally opt to use alternate frame rendering (AFR) for multi-GPU modes as every other frame is rendered by the second GPU. But AFR can introduce latency and create a higher CPU load, which limits scaling. Nvidia has abandoned support for 3-way and 4-way SLI for this reason. But what if there were a better way to scale?

AMD's Instinct MI250X GPUs and Intel's Xe-HPC solutions are tackling the problem, but both are intended for data center workloads rather than real-time game rendering. However, there's enough communication between the GPUs that they have much higher bandwidth interfaces, plus they stack multiple chips to boost performance.

Look at what AMD did with its Ryzen 3000 and Ryzen 5000 CPUs. AMD created the Infinity Fabric as a high-speed link between the CPU chiplet and



AMD's Instinct MI250X data center GPU uses stacked chiplets.

the main chipset, and on higher core-count models there can be anywhere from two to eight CPU chiplets. Instead of separate chips linked by a bus, MI200 GPUs will stack the chips on an elevated fanout bridge with an ultra-high bandwidth die interconnect. Intel's Xe-HPC (Ponte Vecchio) stacks 47 chiplets, linked by a high-speed EMIB (embedded multi-die interconnect bridge). Either way, the goal is the same: increase chip-to-chip bandwidth and enable better scaling across GPUs.

It's like the next evolution of HBM (high-bandwidth memory), where multiple stacks are placed on a silicon interposer that links with the main GPU to provide gobs of memory bandwidth. Except now, instead of memory, Intel and AMD are stacking processors, cache, and whatever else might be useful.

The specs of AMD's MI250X look impressive. Each GPU chiplet has 110 compute units (CUs), 440 matrix cores, and 64GB of HBM2e memory. Combined, the GPUs consist of 58bn transistors, 128GB of HBM2e with 3.2TB/s of bandwidth, and 220 CUs that can deliver 48 TFLOPS of FP64 compute—double that for FP64 work done on the matrix cores. On paper, it's nearly five times the compute of Nvidia's A100 solution and will be deployed in the Frontier and El Capitan supercomputers.

When will we see the consumer equivalent? We don't need 128GB of HBM2e or all that FP64 number crunching, just 120 CUs focused on FP32 graphics performance, linked by a similar high-speed interconnect. Even a single 120 CU AMD GPU made on 5nm technology could be up to 50 percent faster than the current Radeon RX 6900 XT. Imagine linking two or four of those GPUs together in a single monster gaming solution. That's a while off, but I look forward to testing the Radeon RX 8950 XT dual-chiplet card in a few years. Fingers crossed.

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

All we need is 120 CUs of FP32 graphics performance, linked by a high-speed interconnect.

THE LIST

THE TOP SIX VR HEADSETS



6

PLAYSTATION VR \$299.99, www.playstation.com

This PlayStation 4 headset brought VR to a wider audience and, though a little outdated, still offers an experience that's easier to manage than a typical PC rig. It is also compatible with the PS5 through a camera adapter. Sony is working on a new version of PlayStation VR for its latest console, which should address the outdated feel of this device.



5

HTC VIVE PRO 2 \$1,399.00, www.vive.com

The most premium VR headset experience on PC comes with a hefty price tag. It has a 5K display paired with a 120Hz refresh rate and the outcome is gorgeous. It is supremely comfortable and the sense of immersion is peerless. Sure, we'd love it to be tetherless, but we can't have everything. Unfortunately, the controllers let it down a little too.



4

OCULUS QUEST 2 \$299.00, www.oculus.com

This standalone headset offers more freedom than tethered rigs, though you can also pair it to your PC with an additional cable if you prefer. Like Vive, Oculus has been around since the start of VR's emergence over the past decade. This headset features ergonomic halo controllers and a decent array of title releases, making the Oculus Quest 2 an excellent alternative to a fully-fledged PC setup. As one of the most accessible headsets on the market right now, it's a great starting point into VR.

3

HTC VIVE COSMOS ELITE \$899.00, www.vive.com

As one of the forerunners of modern VR gaming headsets, Vive knows what it is doing. The Cosmos Elite is a solid PC VR headset that gives you a great immersive experience. It is well constructed with accurate tracking and the gameplay is as smooth as you'd expect from HTC Vive. Expensive, but the premium feel and the great controllers make it worth the money.



2

HP REVERB G2 \$449.00, www.hp.com

HP isn't particularly well-known for its VR gear, but that isn't reflected in the quality of the Reverb G2. Developed alongside Microsoft and Valve, this PC-based tethered headset provides great immersion. Its high-quality screen at 2160 x 2160 per eye and a decent level of comfort deliver a strong VR gaming experience, all for a fairly reasonable price tag.

1

VALVE INDEX \$999.00, www.store.steampowered.com

A VR headset from the creators of Steam should be a recipe for success. It's fairly pricey and doesn't have the luxury of being untethered, however, its 120Hz display should help prevent motion sickness, one of VR's biggest downsides. The individual finger-tracking controllers are revolutionary, making the Index a great high-end VR package.





Jeremy Laird

TRADE CHAT

8K screens and the future of high-DPI computing

MICROSOFT HAS BEEN PUSHING the idea of high pixel density computing since the release of Windows 7 10 years ago. That it hasn't happened yet—a fact we were reminded of by a recent question on 8K gaming from a reader—is odd, because consumers are used to high-DPI displays on their smartphones.

That smartphone experience generates a sense of cognitive dissonance. Any handset launching with less than 300dpi would get laughed off the internet these days. But most PC users are lucky if their monitor manages 120dpi. Of course, smartphone screens are normally held closer than the typical viewing distance for a PC monitor, but not so as to make up for the chasm between the pixel densities.

Many smartphones have pixel densities of 400dpi or more, with some high-end models reaching 500dpi. Meanwhile, a typical 1440p 27-inch PC monitor manages just 109dpi. Even a 27-inch 4K monitor is only 163dpi. So why does this matter? It's not about content, at least not in the sense of video or gaming. Virtually no content exists that fully leverages the capability of a high-DPI PC monitor. You'd need beyond-4K resolutions, which presents numerous problems. For starters, the 8K format remains nascent. Besides a few pioneering broadcasts, such as the recent Tokyo Olympic games, virtually no TV is made in 8K.

Admittedly, a few movies are now shot in 6K and above, but they are still typically mastered in 4K or below. Special effects are rarely rendered in 6K or 8K, even when that format is used for the primary digital video capture. As for gaming, 4K presents a big enough challenge for today's GPUs, let alone

8K, which contains four times as many pixels as 4K, no fewer than 33 million pixels in fact. Compared with 1080p, 8K has 16 times as many pixels. That's 16 times the rendering load for a graphics card. It's a nightmare for any GPU.

So, the benefit of a high-DPI display isn't anything to do with movies or games, at least not anytime soon. Instead, it's mostly about fonts. Text is much crisper and clearer on a high-DPI display. You get more desktop real estate or at least the option to create more depending on your preferences. That may not sound like a big deal, but I'm surprised that smartphone-savvy consumers are willing to put up with the fuzzy text on PC monitors.

Question is, what level of pixel density is required to deliver a true high-DPI experience? That's a subjective call, albeit one based on empirical fact. It all hinges on the pixel density required to prevent the human eye from resolving individual pixels. Beyond that point, adding further pixels makes no difference. That, in turn, hinges on both pixel density and viewing distances. I'd say 200dpi is a pretty good rule of thumb. So, a 4K 27-inch monitor isn't a million miles



Previous-gen 8K TVs, such as the LG Nano 95, now sell for around \$1,000.

away. But it isn't quite there.

This brings us back to that 8K reader question. As ever in the PC monitor market, trends in TVs will drive prices down. It took a million cheap TVs before 4K monitors went mainstream. The same applies to 8K and thus affordable high-DPI computing.

The prospects look good. The TV industry is pressing on with 8K panels, despite the lack of content. It's already possible to pick up an 8K TV for around \$1,000 if you shop hard. The cheapest 8K monitors still cost thousands more.

Right now, the smallest 8K TV is 55-inches, which somewhat undermines the high-DPI remit. But a 40-inch set would deliver 220dpi. These days, they're pretty much giving away 40-inch 4K TVs. Here's hoping, a few years down the road, the same will be true of 40-inch 8K sets and I can finally have the fonts of my fantasies.

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

The benefit of a high-DPI display isn't really to do with movies or games, it's mostly about fonts.

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

- > Monitor misery
- > Start-up sound
- > HDR headache

DVI to USB-C

Long-time subscriber for more than 20 years here: I have an LG 29-inch 2K monitor, model number LG 29UM59-A. This monitor has one HDMI and one USB-C port. I have been trying to connect a second computer to the USB-C port without any success.

The motherboard with the USB-C port is an MSI X370 Gaming Pro (with AMD Ryzen 1700X CPU and 32GB RAM), with an MSI GeForce GTX 750TI video card with 2GB of memory. I'm hoping to upgrade to an RTX 3070 or RTX 3080 when prices come back down to normal.

I've tried both USB-C-to-VGA and USB-C-to-DVI adapters and neither works. I've not tried a direct USB-C cable yet as I do not have one. Is there some reason why these adapters don't work? Is USB-C bidirectional? **—Roger**

THE DOCTOR RESPONDS:

First up, you won't be able to connect the USB-C port on your PC to your monitor – that's because the USB-C port on your monitor's connector is more accurately described as a DisplayPort or Thunderbolt over USB-C one,

so it needs connecting to your graphics card as you've tried using the two adapters.

The issue with the adapters is that they're one-way, so are designed to connect a DisplayPort/Thunderbolt over USB-C port on your PC to a DVI or VGA port on your monitor. You want something that works in the opposite direction—a DVI to DisplayPort adapter. We found a StarTech.com DVI to DisplayPort adapter for \$39 from CDW (www.cdw.com/product/startech/4888592). You will need to pair that with a bi-directional USB-C-to-DisplayPort cable to complete the connection to your display—Amazon Basics offers a 3ft cable for around \$15 (www.amazon.com/dp/B081VK7Q94). These two items combined should enable you to use your PC with your LG monitor and when you come to upgrade to an RTX 3070 or 3080, you'll be able to use the cable to connect it directly to your monitor, allowing you to retire or redeploy the DVI-to-DisplayPort adapter elsewhere.

Change start-up sound

My issue is more an annoyance than anything

else: I've successfully installed Windows 11 on both desktop and laptop, and on the whole, I'm happy with it. However, I frequently shut down my PC, and rely on being able to hear the Windows logon sound to know it's ready for me to use.

This wasn't a problem in Windows 10, but the Windows 11 start-up sound is too wimpy. Unfortunately, I'm struggling to find a way to change this sound to something better – an online tip suggesting a Registry tweak failed to work. Can you tell me how to change the sound associated with Windows 11 starting up?

—Gerry Exstein

THE DOCTOR RESPONDS:

The Registry tweak Gerry refers to, namely to navigate to HKEY_CURRENT_USER\AppEvents\EventLabels\WindowsLogon and change the 'ExcludeFromCPL' value to 0, allows users to change the Windows Logon sound via the Sounds tab of the old Sound Control Panel (access this via 'Settings → System → Sound → More sound settings').

However, the reason the option is hidden is that the sound in question is now

redundant – it's not the source of the sound you'll hear when you log into Windows.

The actual fix for Gerry's problem lies elsewhere and would normally involve patching a system file (C:\Windows\System32\imageres.dll). Thankfully, the excellent portable (and free!) Winaero Tweaker (<https://winaerotweaker.com>) tool does the heavy lifting for you.

First, you need to source a WAV file of the start-up sound you wish to use—you can download the sounds from earlier versions of Windows from WinHistory.De (www.winhistory.de/more/winstart/winstart_en.htm), for example.

Once downloaded, open Winaero Tweaker and navigate to Appearance → Startup Sound. Clicking 'Play current sound' will confirm the wimpy sound you're hoping to change, then click 'Replace startup sound' to select the replacement WAV file. You'll be prompted to reboot your system, but simply opt to do it later, then click 'Play current sound' to confirm the sound has been changed. Should you later want to switch back, simply click 'Restore default sound'.

∨ submit your questions to: doctor@maximumpc.com

MBR to GPT revisited

After reading your reply to Kirk Holmes in the November issue (MPC 196) about converting from MBR to GPT to support Windows 11, I opened my free copy of Partition Wizard 8 and got some blurb about a new version 12 available. I uninstalled my version 8 and installed version 12. When I select Convert MBR Disk to GPT Disk, I'm taken to an online page that requires me to pay for one of MiniTool's pro subscriptions to use the conversion tool.

I don't want to pay for a tool that I probably will never use again. If I take the MBR2GPT approach, what happens to my programs and data? Will Windows have to be reinstalled? Alternatively, if I create a disk image using my free version of Macrium Reflect, then convert the SSD from MBR to GPT, will Macrium be able to restore the disk image to the new GPT partition with no issues? Can you suggest a better way to get Secure Boot properly enabled? I sent this question to Macrium support, but the only response I got is that they don't provide support for the free versions of Reflect.

—Steve Vanetti

THE DOCTOR RESPONDS:

Apologies to you and anyone else who discovered that MiniTool (along with other partition managers) have conveniently removed the non-destructive MBR-to-GPT conversion tool for system disks from their free versions. Thankfully, Microsoft has thought to provide a non-destructive path via its recovery tools.

To answer your question about Macrium, you can use it to image the drive, convert it to GPT using a destructive method and then restore the image to the new GPT partition—why Macrium's engineers didn't direct you to their own knowledgebase escapes the Doc. Now version 8 has been released for free users, head over to <https://>

knowledgebase.macrium.com/display/KNOW80/ and navigate to 'Restoring and Browsing → Restoring an MBR System image to UEFI/GPT' in the left-hand pane for full instructions.

If you'd prefer to follow Microsoft's non-destructive procedure, the Doc advises that you take a drive image backup (just in case) and make sure you have Macrium rescue media handy too. Once done, open Settings to Update & Security → Recovery and click 'Restart now' under 'Advanced startup'. When the menu appears, select Troubleshoot followed by Advanced options and finally Command Prompt. Next, type the following command:



8-bit HDR colors, pictured on the right, are more accurate but look less vibrant on a monitor.

`mbr2gpt /validate`

Assuming no errors are found, type the following: `mbr2gpt /convert /disk:0`

This will convert the first drive in your system (if you have a non-standard drive setup, use the 'diskpart' command followed by 'list disk' to obtain your system disk's number using its size to identify it) to GPT. Once complete, reboot your system and enter your PC's BIOS setup to switch its boot mode from legacy to UEFI.

Once done and saved, you should find you're able to boot back into Windows. At some point in the process, you're likely to get a warning about MBR2GPT failing to update ReAgent.xml—you'll need to manually disable and re-enable the WinRE environment. To do this,

right-click Start and choose Command Prompt (Admin), then type the following:

```
reagentc /disable
reagentc /enable
```

This should fix the WinRE error, enabling you to enable Secure Boot and upgrade to Windows 11.

HDR? Meh!

I made my first foray into HDR with a Gigabyte G32QC-A 1440p HDR monitor. I had to go with 1440p again since I'm stuck rocking a GTX 1080Ti. I'm running it on a Ryzen 5900X/X570 & 64GB PC3600 DDR4 Windows 11 full release system, so HDR should be fine.


I hooked up the G32QC-A via DisplayPort 1.2, installed

in the NVIDIA Control Panel and override enabled on the monitor.

Did I get a bum monitor, or is this an aspect of HDR nobody talks about, except all the other end users reporting the same thing with various HDR monitors? Your thoughts are appreciated. —BJ Koho


THE DOCTOR RESPONDS: This appears to be a limitation of the monitor itself. It's certified as DisplayHDR 400, described as a "first genuine entry point for HDR" (<https://displayhdr.org/#tab-400>). The main issue is that, unlike higher-end HDR displays, the G32QC-A only supports 8-bit color. The blurb states it "responds to HDR10 content,

putting its generous color gamut to use and allowing a 10-bit color signal to be used for enhanced precision". What this means in reality, is that the G32QC-A is unable to display all colors from the 10-bit range on an 8-bit display. It chooses a close match but will lack the vibrancy of a true 10-bit HDR monitor because it relies on GPU dithering to display more accurate colors.

Finally, regarding your problems getting the monitor to work at 165Hz, you might want to check which display resolution you've set via the NVIDIA Control Panel—make sure you select 1920x1080 under the PC section rather than the default 'Ultra HD, HD, SD' section, and you should see the 165Hz option is available to select. 

I have spent hours going through every setting and tried tweaking everything, but SDR has way more pop and is clearer than HDR. I also can't get it to offer any refresh rate higher than 120hz (165hz monitor), even though the monitor supports NVIDIA cards through adaptive sync, and I have 'G-SYNC compatible' enabled

OSD Sidekick, made sure I had the latest NVIDIA drivers and monitor firmware, followed every online tutorial I could find. Then I jumped into *Cyberpunk 2077* in full-screen mode and hit Win+Alt+B. But instead of blow-my-socks-off color, the screen dimmed slightly with muted, if more accurate-looking, colors and everything looked like I was viewing it through a thin fog.



2022 TECH PREVIEW

THE *MAXIMUM PC* TEAM LOOKS INTO ITS CRYSTAL BALL TO GIVE YOU A HINT OF WHAT'S TO COME IN 2022

PANDEMIC ASIDE, 2021 has been one heck of a year. Our world of technology is still forever advancing, regardless of the silicon shortages. We've seen new graphics cards, processors, screens, peripherals, cases, cooling solutions, and ever crazier announcements and launches over the past 12 months. The tech landscape has really changed quite dramatically, companies have thrived, died, and survived as various economic impacts of the silicon shortage have finally hit home. And yet, the long advancing march of technology still trundles on ever forwards.

So, 2022 is set up to be quite the year. Hopefully, its arrival will signal not only the end of the pandemic and its many impacts but also the end of the chip shortages too. With a host of new fabs and technologies opening up across the globe, supply chains re-fortifying, and crypto mining on the decline on anything that isn't an ASIC, you might finally be able to buy a graphics card at a reasonable price (as long as the scalpers don't get there first), or at least by the time you're reading Tech Preview 2023 anyway.

So then, with the past year out of the way, it's time for us fine folk to give you our best predictions for the next year. We've already seen the launch of gaming changing new technology with the likes of Intel's hybrid Alder Lake parts, JEDEC's DDR5 RAM, and more, but what else does the future hold for tech-heads? Well, we got some of the best people in the business to give you the inside scoop on what we might see going forward into the next year and beyond.

© NONSITZ/BETTY IMAGES

CPUS & GPUS

CPUS: ZEN 4 TAKES ON RAPTOR LAKE IT'S INTEL VERSUS AMD ALL OVER AGAIN

2022 IS SHAPING UP to be potentially one of the more exciting years in the CPU realm. After AMD's Zen 3 CPUs claimed the overall performance crown in 2020, Intel responded in late 2021 by finally releasing a desktop CPU built on something better than 14nm—Alder Lake used the Intel 7 process, aka the artist formerly known as Enhanced 10nm SuperFIN. For 2022, we'll be getting two more contestants for the CPU battlegrounds: AMD Zen 4 and Intel Raptor Lake.

Zen 4 represents the next big jump for AMD and will be the first Ryzen desktop solution to move on from socket AM4. Zen 4 will move to a new AM5 socket, with plenty of features that required a change in sockets. Will AM5 also ditch the CPU pins and move to an LGA design? All signs indicate that will happen. Zen 4 will target general IPC and feature improvements, and will also be the first x86 CPU to move to TSMC's N5 5nm process technology.

The move from 7nm to 5nm promises the usual gamut of chip size, power, and transistor density improvements. N5 entered volume production in 2020 with Apple's A14 Bionic processors, so it's already had plenty of time to mature. TSMC says N5 can improve transistor density by about 1.8x over N7, with 30 percent lower power use at the same performance, or 15 percent higher performance at the same power. More importantly, while the recently rechristened Intel 7 process might be competitive with TSMC's N7, it's still a generation behind N5, and AMD plans to leverage that lead.

AMD hasn't given any specific statements regarding performance and architectural improvements with Zen 4, but we do know the consumer Raphael CPUs, which will be sold as the Ryzen 6000 series, should arrive before the end of 2022. AMD is likely to stick with a maximum 16-core CPU design for the consumer segment, which leaves plenty of extra transistors for other features. Note that AMD has also revealed plans for a 96-core Genoa Zen 3 data center part in 2022, with a 128-core Bergamo Zen 4 chip planned for 2023.

Zen 4 will also join the DDR5 bandwagon. Officially, AMD hasn't said whether Zen 4 will require DDR5, so it could end up supporting DDR4 and DDR5—though, like Alder Lake, motherboards would almost certainly use either one or the other

memory type. PCIe Gen5 support should also be present, though it may not be enabled on all motherboards or chipsets. Power and motherboard cost will be driving factors in both memory and PCIe support.

Early claims suggest that Zen 4 may deliver up to a 25 percent IPC (instructions per cycle) improvement over Zen 3—or maybe just 25 percent more performance. Either would be great, as even 10–20 percent gains are hard to come by. More cache, improved latencies, and other architectural enhancements might be able to get there. Zen 4 may also see Ryzen CPU clocks break the 5GHz barrier for the first time.

Intel's Raptor Lake will be less revolutionary, relying instead on modest changes relative to the recently launched 12th-Generation Alder Lake CPUs. Raptor Lake will be 13th-Generation Core and should maintain socket compatibility with LGA1700. That means it will continue to support DDR4 and DDR5, as well as PCIe Gen5.

The big change with Raptor Lake, other than the usual incremental generational improvements, appears to be in cache sizes and core counts. Both are based on early leaks and rumors, but we may see much larger L3 caches on Raptor Lake, to help close the gap with AMD's Zen 3 and V-cache models.

As for the CPU core counts, Alder Lake has a maximum of eight Golden Cove performance cores with Hyper-Threading and eight Gracemont efficiency cores that only run a single thread each. That gives a maximum 16-core/24-thread configuration on the Core i9-12900K. Raptor Lake will improve on that number, but interestingly the current word is that the maximum configuration will be eight Golden Cove cores (16 threads) and 16 Gracemont cores (16 threads as well), yielding a maximum 24-core/32-thread design. Raptor Lake is also reported to target efficiency more than pure performance, which would coincide with the addition of more efficiency cores while keeping the same number of performance cores.

Intel has only recently launched its Alder Lake CPUs, but the 13th-Generation Raptor Lake should use the same socket and motherboards and further improve performance and efficiency

TO GPU OR NOT TO GPU?

Intel has been including integrated graphics solutions with all of its consumer designs since the 2nd Gen Sandy Bridge family launched back in early 2011. More recently, it has begun selling F-series chips that lack graphics, but those are just the same CPU die with the graphics fused off. In contrast, AMD has so far split its Ryzen family into CPUs that lack graphics and APUs that include them.

With Zen 4, AMD may stop bifurcating its processors. The business and OEM sectors love all-in-one chip solutions, particularly in the growing laptop market. With 80 percent more transistors in the same die area, even with other architectural improvements, it wouldn't be a stretch for AMD to include some form of GPU in all Zen 4 chips. Unfortunately, RDNA 2 graphics tech may be used, it will likely be significantly underpowered relative to dedicated graphics cards.



GPUS: IT'S LOVELACE, RDNA 3 & ALCHEMY, OH MY!

2022 TRIPLE WHAMMY

IT'S NO SECRET that the world of PC graphics has had problems delivering sufficient quantities of cards ever since the current generation Ampere RTX 30-series and RDNA 2 RX 6000-series parts launched in late 2020. Nevertheless, both AMD and Nvidia should be forging ahead with their next-generation RDNA 3 and Lovelace architectures in 2022, and they will be joined by Intel's Arc Alchemist. Will you be able to buy a card? Let's hope the miners find somewhere else to dig.

Officially, neither company has said much of note about their next-gen GPU architectures. We expect AMD to move from TSMC N7, which it has used on the past two generations of RDNA and RDNA 2 GPUs, to TSMC N5, keeping the latest generation CPUs and GPUs on the same node.

As with CPUs, N5 should enable a lot more transistors in the same area, along with better power and performance. RDNA 2 saw some impressive gains from the large Infinity Cache, and there's no reason to expect RDNA 3 won't keep that big cache. We could see a move to a stacked die holding the cache, but that's potentially more expensive and with the larger transistor budget likely to be unnecessary.

Navi 21 currently has a modest (for a halo GPU) die size of 520mm square, with nearly 27 billion transistors. RDNA 3 and Navi 31 (or whatever it eventually gets named) could keep the same die size and increase to almost 50 billion transistors. A 384-bit memory interface and 120 CUs seem possible. Architectural improvements to increase ray tracing and general shader performance should yield a GPU that's up to 50 percent faster than the RX 6900 XT. It's also possible AMD could go with a multi-chip approach and with that a truly massive bump in performance (see boxout, above right).

Nvidia meanwhile has something to prove with Lovelace. It's not that Ampere isn't good—sometimes, it's too good when it comes to mining performance—but efficiency took a big step back compared to previous designs. Also, the reference model

AMD's Zen architecture should march forward in 2022, with the next generation Zen 4 potentially delivering up to 25 percent more performance

MONOLITHIC DIES VS. CHIPLETS

Most graphics cards have traditionally used monolithic dies—everything built into a single chip. It makes signal and data routing faster, but it does result in large chips. AMD may choose to go a different route with RDNA 3, and Nvidia and Intel may do the same with future GPUs.

With planar scaling of semiconductors slowing down, opening up the third dimension with stacked dies and multiple chiplets could be the solution. Making smaller chips and linking them together also reduces the cost and the time it takes to get to market.

There are rumors that AMD will use its proven chiplet approach from CPUs with RDNA 3, with smaller chiplets interlinked by the Infinity Fabric. Done properly, this could allow AMD to scale from budget to extreme GPUs without the need to make three or more different chip designs.

We also know AMD is using a chiplet approach with its MI250X data center solution, and Intel's data center Xe HPC uses multiple 'tiles' as well, plus 3D chip stacking. We may not see those used in RDNA 3 or Alchemist, but don't be surprised if the next-generation products from all three companies opt for multi-chiplet designs.

Founders Edition cards can run incredibly hot thanks to the GDDR6X memory. Both deficiencies should be addressed with the RTX 40-series Lovelace cards.

Lovelace is also expected to move to TSMC's N5 process, joining AMD and Apple on the cutting edge. Nvidia perhaps got a bit complacent with Ampere, figuring that Samsung's 8N could stay competitive with TSMC's N7. Supply and prices may have been lower on Samsung Foundry as well, but going from a 260W RTX 2080 Ti to a 350W RTX 3080 Ti basically scaled power use with performance. Jumping straight to N5 should bring drastic reductions in power.

It's not just power, of course. 5nm should provide for substantial increases in transistor density and performance. GA102 was a large 628mm square die—20 percent larger than AMD's Navi 21, despite only having a billion more transistors. Some of that is due to Infinity Cache, but we expect Nvidia to go after a 600mm square die on N5, which would mean over 50 billion transistors.



How many GPU cores would that allow? Some rumors are suggesting up to 144 SMs and 18,432 CUDA cores on the top model, and that would likely be a harvested chip. Combine that with fourth-generation tensor cores, third-generation RT cores, more cores in general, architectural enhancements, plus potentially higher clocks and more cache, and Lovelace should be able to keep Nvidia in the pole position.

But it's not just AMD and Nvidia this time. Intel's Arc Alchemist (aka DG2) cards should finally launch in early 2022. We've talked a bit about Arc over the past year, and there are plenty of unknowns. Still, with the top cards reputed to contain 512 vector engines (the new name for Intel's Execution Units of previous integrated GPUs), or 4096 GPU cores running at up to 2.3GHz, plus lots of GDDR6 memory, Intel Arc should perform decently.


Let's be real: It's not going to match Lovelace or RDNA 3. But it might be able to compete with the existing RDNA 2 and Ampere products. More importantly, it may be available at a reasonable price, with the top model reportedly selling for \$650 (according to Intel's own accidental leak). If it can keep up with an RTX 3070 and actually have stock at that price, that would be a great starting point. The sequel to Alchemist, Battlemage, can then worry about closing the gap. Let's just hope drivers are up to snuff.

CASES & COOLING

RETHINKING HOW WE CHILL NANO? NO, NO

LAST YEAR WE SPOKE in-depth about how the world of cooling PCs was changing exponentially quickly. We'd seen the likes of nano-particle coolant enter the fray, courtesy of Mayhems, and were convinced it was going to make its way into the world of AIOs. Sadly this just hasn't come to fruition, and as the genius behind Mayhems has stepped back from the world of development, we'll be curious to see if any of that innovation makes its way into the world of AIOs. That said, the world of CPUs still hasn't changed. New generation chips are still as hot as they've ever been, and although we've not seen a new generation of processor from AMD just yet, it's highly unlikely we'll see a reduction in heat.

AMD's Instinct MI250X uses dual GPU chiplets and stacked dies to achieve record performance. Does this represent the shape of things to come in 2022?



© AMD, MAYHEMS, CORSAIR

The answer to these heightened temperatures so far seems to come from larger and larger radiators. Corsair launched an industry-first 420mm AIO early in 2021, and it's something that we're likely to see start popping up from other manufacturers too. In the short term, it certainly makes sense. It's far easier to just bump up radiator sizes than mess around with unpredictable new coolant, especially for mass-produced products, and although nano-coolant showed potential, it wasn't without its flaws—namely, cost and production time.

Given many case manufacturers also create their own coolers, we'd be surprised not to see some 420mm AIOs pop up from the likes of Cooler Master. Expect to see more AIOs with in-built displays, and potentially even a return look at expandable AIOs, compatible with graphics card water blocks and beyond. We've already seen Asus co-design a GPU with Noctua this year, extending that notion to graphics cards, and making it compatible with a liquid-cooled loop would be the logical next step.

CASE SHIFT

ANOTHER UNIQUE WAY of combating the problem with heat comes from case design. With the success of more open-style air-flow cases, from the likes of Hydra and Yuel Beast, we'd put money on some of the more mainstream case manufacturers following suit.

You can already see a lot of that design language going into Fractal Design's Torrent chassis, with its broad focus on airflow, and vast gaps on its intake fans. This will inevitably lead to larger build-ups of dust, but it does allow much better access to air. We even cooled a Core i5-10600K using a slim-line Noctua NH-L9i without issue, purely by using an open-air chassis.

Case design takes time, and bringing a product into production typically takes anywhere between 12 to 24 months. It may only be a matter of time before Fractal, Phanteks, and NZXT launch something a bit bolder than the current airflow variants of their more famous Evolv, H510, and Define lineup of cases.

EOL FOR MICRO-ATX & DRIVE BAYS

SADLY, this might also be the final year that we see the availability of two specific types of cases—at least, at the premium end of the market. With more and more high-frequency chipset launches and a dwindling fan base due to GPUs moving away from Crossfire and SLI support, Micro-ATX seems to be dying a slow death at the hands of case manufacturers.

While available to the public at a more reasonable price, sadly nano-coolant hasn't made its way into AIOs just yet



It's a typical economic cycle, there's less demand for M-ATX cases, so fewer are made and it becomes more expensive to produce, and so fewer are bought, and so on. With most of those looking for compact cases aiming for ITX instead of the Micro form factor, this could be the year that those cases built with Micro-ATX support finally bite the bullet, and it turns into a footnote on the specs of most ATX tower motherboard support instead.

Another beloved spec lined up for the chop is the humble 5.25-inch drive bay. As manufacturers are following the trends of PC gamers, the humble disk drive bay may not be long for this world. We looked at the market last issue, to find out how many premium 5.25-inch drive bay supporting cases are still available. Our list came to just six, most of which were from be quiet! or Fractal Design, and the vast majority were more than two years old.

There's a chance the Swedish and German companies might be the last bastions of the standard, but with many mainstream manufacturers already ditching the 5.25-inch slot in favor of a more compact design, and better cooling capacity, if you're set on keeping hold of these plucky drive bays, we recommend you grab a case that supports one sooner rather than later.

STORAGE

SOLID STATE OF THE NATION

GET READY FOR PCI EXPRESS 5.0

2021 WAS THE YEAR Intel belatedly spluttered onto the scene with support for PCI Express Generation 4 via its Rocket Lake CPUs. That came nearly two years after AMD debuted PCIe 4.0 support with its Ryzen 3000 chips. So, could Intel be left behind yet again in 2021 with the arrival of PCIe Gen 5?

First, let's cover off what the impact of PCIe Gen 5 will be on the broader solid-state storage landscape next year given that the PCI Express interface dictates the maximum theoretical speeds of solid-state drives aimed at consumers. Japanese memory specialist Kioxia has been among the early movers in Gen 5 drives. It has shown off a prototype drive that's claimed to be good for an epic 14GB/s sequential read performance.

That's not far off the circa 16GB/s theoretical limit for a quad-channel PCIe 5.0 drive. Write speeds for the Kioxia drive are admittedly a little less stellar at 7GB/s.

Corsair brought the first 420mm AIO to market, but it won't be the last

Hydra's open-air design may become more common in future



FAN...TASTIC NEWS? NOT REALLY

Fan development has come to a bit of a standstill recently, with little development from some of the world's best fan makers. The 120mm and 140mm fan sizes are still likely to be the most prominent going forward, but aside from that, there's nothing major to report here.

Noctua recently debuted chromax black variants of its latest NF fan designs (the company seems to have finally cottoned on that black fans might be more popular than that traditional beige and brown combo), but as far as development goes, the main aim for most seems to be designing better and more potent RGB fans.

RGB's need for suitable opaque plastics has often led to the plucky blowers producing less than adequate airflow and static pressure, as the plastic is typically flimsier in contrast to the solid fan designs produced by the likes of Noctua or Corsair's ML series. This year, we expect to see a new generation of RGB fans, exploring newer, more durable, yet reflective materials. So there will be some change, albeit minor.

Noctua's passive cooler was a big part of our coverage last year, and it was great to see this come to market. But it has taken the company some time to bring it to fruition and mass production. It's not entirely perfect either, not being rated for some of the more potent AMD processors. We don't expect other companies to suddenly jump on the passive cooling bandwagon, but we'll likely get a 1.2 revision of the P1 capable of chilling whatever AMD decides to launch this year, all the way up to the top.



techradar.



The home of technology

techradar.com

But that's still quicker than any PCIe 4.0 SSD. Of course, peak sequential performance isn't the whole story. Actually, it can be a distraction if it means random access and latency are ignored. But the prototype drive looks strong across the board.

Peak IOPS are quoted at no less than 2.5 million (the best PCIe 4.0 drives hit around one million IOPS), while latencies of 70µs read and 10µs write are mooted. In other words, drives based on the PCIe 5.0 interface promise to be a big step forward by most metrics. This brings us back to the other critical question, namely, platform support.

Intriguingly, it looks as though the tables could be turned for Gen 5. By that we mean it might well be Intel that delivers support first, on the desktop at least. The following does not apply to enterprise platforms. Anyway, at the time of writing in late 2021, Intel has only just launched its Alder Lake CPUs, complete with a new Z690 chipset and PCIe 5 support. Meanwhile, the latest rumors concerning AMD's next-gen Zen 4 based CPUs suggest that PCIe 5 isn't quite guaranteed to be part of the feature set.

If that looks like a total turnaround in favor of Intel, there's a catch and it concerns the way PCIe support is partitioned on Intel's Alder Lake platform. Leaked chipset diagrams indicate that PCIe Gen 5 support on Alder Lake will be limited to the 16 primary CPU lanes—in other words, those typically populated by a graphics card on a performance PC. The storage-specific quad-lane interface that also connects directly to the Alder Lake CPU remains Gen 4 spec.

It will still be possible to use a Gen 5 SSD with Alder Lake. But that will require splitting the 16-lane graphics interface in two, thereby limiting the graphics card to eight lanes. In reality, that's probably plenty for any conceivable graphics card. Eight lanes of PCIe 5.0 is equal to 16 lanes of PCIe 4.0 and nobody claims 16 PCIe 4.0 is insufficient for modern graphics. Arguably more problematic is that you'll probably need a PCIe-16 adapter card to run a Gen 5 SSD. The chipset partitioning means the Z690's M.2 SSD slots will be limited to Gen 4 operation.

Cases that support 5.25-inch drives could become increasingly rare, so grab one while you still can



© BE QUIET!

Elsewhere, we're expecting SSD progress in 2022 to be largely incremental. When it comes to memory density, a number of flash memory chip producers have recently introduced 176-layer TLC NAND technology with bit densities up to 10Gb/mm², while Intel's latest 144-layer QLC chips weigh in at just under 14Gb/mm². We're not expecting any advances on that in 2022. Instead, you can expect to see prices for SSDs of all types and capacities come down, thanks to wider adoption of the more dense NAND chip tech that appeared in 2021 and a broader relaxation of the global chip shortages we've seen in the last 18 months or so.

As for the rest, don't expect to see any revolutions. PLC or penta-level NAND flash is still years away. Similarly, Intel's new second-gen 3D Xpoint chips are still miles off the data density of NAND memory, so its Optane drives won't come close to the capacity of NAND-based SSDs for anything like similar money.

DISPLAYS

MINILED GOES MAINSTREAM

HOMING IN ON TRUE HDR

WILL 2022 be the year that HDR becomes an actual thing for PC monitors? That will probably hinge on the impact of miniLED backlighting, both in terms of the effectiveness of the technology itself and the extent to which it is adopted.

Our first taste of miniLED was inauspicious, to say the least. Samsung's so-called 'Neo' G9 monitor promises much on paper. Its miniLED backlight packs no fewer than 2,048 local dimming zones and Samsung claims a peak brightness of 2,000 nits. At last, we're getting HDR instead of pseudo-HDR monitors with a handful of edge-lit zones or at best a few hundred.

Or so we thought. The Neo G9 was plagued with issues, some of which were the result of sloppy product development. A firmware update addressed some of the issues, most obviously it allows the Neo G9 to get close to that claimed 2,000 nit peak where the original firmware meant you were lucky to see half of that.

NEWCOMERS TO THE ARENA

We've also had a few newcomers rear their heads up over the past 12 years, namely Razer. The company has gone from strength to strength and has just announced a new product to its lineup, a selection of AIO liquid-coolers and fans, ranging in sizes from 240mm to 360mm.

Combine that with a new generation of Tomahawk cases and we expect Razer will push into this market. Asus and MSI, both of which have enjoyed significant success recently, are also likely to hit the ground running this year with more case launches.

Whether these will be fleeting jaunts into the case and cooling market or more permanent ventures is anyone's guess. The cooling and case market is notoriously difficult for carving out a niche, however, given the current state of the chip industry and with manufacturers trying to widen their revenue lines, we'd be surprised if more didn't follow suit.

We're even hearing rumors that HyperX might be investing significant assets into building its own case lineup. But for the time being, it's just that, rumors.



Razer may bring some sophisticated designs to market, but will it be enough to carve out a segment of its own?

But the Neo G9 also highlighted the fundamental shortcomings of an LCD monitor with local dimming technology as opposed to display technology with true per-pixel lighting, such as OLED. Long story short, with local dimming, the backlighting is subject to algorithms that will inevitably throw up various artifacts and issues. So long as the lighting is not per pixel but, in fact, shared over a number of pixels, then there will always be compromises. And it's worth recalling that even with 2,048 dimming zones, each zone on the Neo G9 is, in fact, lighting a block of 3,600 pixels.

Put another way, those 2,048 dimming zones work out to a backlighting resolution of around 70 by 30. Imagine if the actual resolution of the display was just 70 by 30 pixels. Right, anyway, miniLED backlighting will be more widely adopted in 2022, but it won't quite be the killer technology it's sometimes cracked up to be. It also won't be cheap. A more reasonable expectation is for wider adoption of OLED.

Admittedly, OLED technology has its own problems, especially when it comes to PC usage, most notably image persistence or burn-in. But the biggest player in OLED panels, LG Display, is both hell-bent on bringing down OLED panel costs and committed to PC-friendly features such as low latency and burn-in mitigation. So, an OLED PC monitor will be that little bit more attainable and practical in 2022. That's especially true with the news that LG plans smaller, cheaper, and more desktop-friendly 42-inch OLED panels for 2022.



Will PCIe 4.0 drives such as the Samsung 980 Pro be left behind in 2022?

Elsewhere, don't expect any dramatic advances. Peak refresh rates will increase to 480Hz, but even with existing 360Hz gaming monitors, the returns compared to 240Hz panels were already diminishing. Of more benefit will be developments in interface technology. HDMI 2.1, for instance, makes 120Hz at 4K possible, something that DisplayPort 1.4 can already achieve. HDMI is more a consumer electronics interface than a pure-PC solution and HDMI 2.1 is more relevant to games consoles than PCs.

However, because screens for consoles are sold in such huge numbers and aimed at a more mainstream market, you can expect to see a new generation of relatively affordable 4K 120Hz HDMI 2.1 panels appear that should also make good PC monitors. In fact, our prediction for the best monitor buy in 2022 is just that. A 32-inch 4K 120Hz panel enabled by HDMI 2.1 and clocking in close to \$500. The snag? You'll need one of those unobtainium current-gen graphics cards to add HDMI 2.1 capability to your PC.

Meanwhile, DisplayPort 2.0 will make both high-refresh 8K computing and 16K resolutions possible. In reality, neither will be much of a thing in 2022. But we'll still look back on 2022 as the year that those resolutions and refresh rates were first possible.

The rest of the monitor market will be a case of incremental upgrades. USB-C connectivity will become more widespread and more affordable and it's also due an upgrade in power delivery from a maximum of 100W to a much more useful 240W. But we're not predicting any big changes in mainstream panel tech. IPS and



Samsung's Neo G9 miniLED monitor disappointed in 2021

VA panels will continue to dominate. They will get a bit faster and a bit brighter. But with any luck, the biggest change will be access to the same screen tech from 2021, just a little cheaper.

SYSTEM MEMORY

THE STATE OF MEMORY IN 2022

THE TRANSITION FROM DDR4 TO DDR5 HAS ONLY JUST BEGUN

DDR5 MEMORY IS at the start of its likely five-year lifespan. This exciting, speedy memory has sure put a smile on our faces. It has been a while now since we had a new version of RAM—just over seven years, in fact. To put that into perspective, it was the same year the Amazon Fire TV Stick was released and the same year that Intel's first 14nm chips were released.

So, it's about time we got an upgraded RAM series. At the start of a new generation of memory, the same questions fly around PC enthusiasts' heads. Should you upgrade to this new memory, should you wait around for prices to drop, and how does it compare to the previous DDR4 RAM? All of these are important things to consider, especially if you are in the process of upgrading your setup or buying a new system. As it's released before the new year, 2022 should see widespread adoption of DDR5 usage, with the following year hopefully making DDR5 the most popular choice of memory, pushing DDR4 off its pedestal.

Like most generational changes, one of the main upgrades is the increase in capacity. When DDR3 was phased out for DDR4, 2GB and 4GB sticks were fairly common configurations. You could get 8GB sticks of DDR3 but this wasn't routine and when DDR4 came in, this capacity became the standard across its lifespan. This wasn't a huge upgrade in capacity, but for next-gen DDR5, that is set to change. On the face of it, the single-DIMM capacity for DDR5 is greatly improved over DDR4. You can expect individual DIMM sticks with up to a whopping 128GB capacity,

and potentially even more in the future. Compared to the 32GB maximum capacity on the consumer market for DDR4, that's four times greater.

Voltage also alters this time around, reducing from 1.2V to 1.1V. Where the real change is found is how DDR5 handles power. With DDR4, the motherboard handled the memory power. This time around, the VRM has moved from the mobo to the stick itself. The memory kit manufacturer now controls the voltage regulations instead of the motherboard company. It simplifies motherboards but makes DDR5 more complex and unsurprisingly more expensive. Some 32GB kits running at DDR5's minimum speeds have been seen for sale for around \$300. That's nearly double the price of similar DDR4 kits, but it's no shock. A new generation of memory is always expensive at launch.

Another big change is of course the speed. The current minimum RAM speed on DDR5 starts at 4800MHz and goes all the way up to 6400MHz. You can get higher than 4800MHz from DDR4, but the bandwidth speeds will certainly be quicker on the newer generation. With core counts increasing on CPUs, these chips require greater memory bandwidth. One architectural change on DDR5 is that it will be split into two sections, each with its own independent 32-bit channels per stick. This now also moves the key in the pins to the center of the DIMM sticks. This means that the standard 2x 64-bit channel DDR4 system will change to a quad 32-bit channel layout with DDR5. Combine this change with greater capacities and faster RAM speeds, you can expect performance almost twice as fast as DDR4, with lower latency speeds and increased bandwidth.

With Intel's Alderlake chips releasing at a similar time and supporting both PCIe 5.0 and DDR5 straight out of the box, AMD's AM5 socket is sure to follow suit with a release sometime in 2022. With Windows 11 also just being released, there are quite a few generational changes all happening at once. This will make 2022 an interesting year and expectations for PC upgrades are promising. As we said earlier, we expect to see a flood of DDR5 DIMM sticks hit the market and gradually become the mainstream choice of RAM. New variants will arrive throughout the year with greater capacities and speeds. If you are about to buy a setup, then, DDR5 will currently feel quite pricey. It could pay to wait a while. But without any doubt, DDR5 the future of memory and will take over from DDR4. It's just a matter of time. 🕒

DDR5 promises up to double the performance of DDR4. As a new tech, it will be expensive in 2022, but will be the future of memory



MAXIMUM PC'S 17TH IRREGULAR

GEEK QUIZ



WELCOME TO OUR ANNUAL
BRAIN-TEASING SESSION *BY CHRIS LLOYD*

HAVE YOU BEEN paying attention this year? Have you been reading *Maximum PC* all the way through? We hope so because it's time to put your PC knowledge to the test. Welcome to our annual geek quiz where we'll see how high your tech knowledge ranks. We've put together seven sections of seven questions each, covering hardware, software, the internet, and the weird side of technology, and including a few random nuggets. So, are your geek credentials bona fide, or are you still a noob? Time to find out.

No taking the easy option—stay away from the search engine, close your browser, and put the phone down. It's time to pit your wits against our cornucopia of conundrums. If you've been reading *Maximum PC* diligently, you should have seen the answers to at least a few of them.

Some of the questions should be easy, some definitely aren't. And no, there is no prize on offer, unless you count the warm glow of satisfaction if you can prove yourself worthy. Math tells us that by simple random guesswork you'll score around 12 points, so let's find out how much better you can do. You may begin.

Processor Posers

1 This year, Intel introduced us to a new acronym for its processors: **MTP**. What does it stand for?

- A] Minimum Thermal Phasing
- B] Minimum Track Profile
- C] Maximum Turbo Power
- D] Maximum Temperature Profile

2 The new Alder Lake chips use high-efficiency Gracemont cores, how much L2 cache do these have?

- A] 2MB per module
- B] 2MB per core
- C] 2GB per module
- D] 3MB per core



3 During Intel's Innovation event, a Core i9-12900K was overclocked to a record speed, how fast?

- A] 6.8GHz
- B] 6.9GHz
- C] 7.0GHz
- D] 7.2GHz

4 Apple's M1 processor has roughly how many transistors?

- A] 13,500,000,000
- B] 16,000,000,000
- C] 18,000,000,000
- D] 19,750,000,000



5 Have a guess at how many versions there are of AMD's Cezanne APUs.

- A] Er, about 18
- B] Around 20
- C] Roughly 23, probably
- D] At least, say, 26



Graphical Gear

8 Which one of these isn't a technology used in anti-aliasing?

- A) MSAA
- B) FXAA
- C) TAA
- D) PXAA

9 Nvidia's GA102-200 GPU from the GeForce RTX 3080 features which configuration?

- A) 336 Tensor cores, 84 RT cores, 10,752 Shading Units, and 112 ROPs
- B) 272 Tensor cores, 48 RT cores, 8,704 Shaders, and 96 ROPs
- C) 272 Tensor cores, 84 RT cores, 10,240 Shading Units, and 96 ROPs
- D) 272 Tensor cores, 80 RT cores, 8,704 Shading Units, and 96 ROPs

10 What is the maximum possible resolution you can get out of any Ampere-powered GeForce card?

- A) 7,680 x 4,320
- B) 8,640 x 5,160
- C) 9,600 x 5,508
- D) 10,560 x 6,240



11 The AMD Advantage specifications for high-end gaming laptops define a minimum fps at 1080p, what is it?

- A) 80fps
- B) 90fps
- C) 100fps
- D) 120fps

12 AMD's Radeon RX 6600 has 8GB of GDDR6, how fast is it?

- A) 224Gb/s
- B) 224GB/s
- C) 256Gb/s
- D) 256GB/s

13 Which graphics API was initially called OpenGL next?


- A) Vulkan
- B) Mantle
- C) Metal
- D) Direct3D

14 Which of these is a codename for a future version of Intel's Xe-based graphics card?

- A) Battlemage
- B) Celestial
- C) Druid
- D) All of these cool names are real

6 AMD will be adding V-cache to Zen 3 chips next year, how much L3 was on the prototype Ryzen 9 5900X?

- A) 96MB
- B) 128MB
- C) 176MB
- D) 192MB



7 Intel still sells Pentiums, there were six new ones this year, so what architecture is inside the Pentium Gold?

- A) Kaby Lake
- B) Comet Lake
- C) Coffee Lake
- D) They're still using Broadwell



Bits and Bytes

15 Intel's 12th gen chips have brought DDR5 support, specifically what?

- A] 1.1V modules up to 128GB with 72-bit data channels
- B] 1.1V modules up to 256GB with dual 40-bit data channels
- C] 1.1V modules up to 128GB with dual 72-bit data channels
- D] 1.2V modules up to 256GB with dual 40-bit channels

16 Seagate has built a dual activator hard drive which can rival some SSDs, how fast do you reckon it is?

- A] A sustained read/write of 524MB/s
- B] A sustained read/write of 584MB/s
- C] A sustained read/write of 614MB/s
- D] A sustained read/write of 644MB/s

17 Intel's new CEO, Pat Gelsinger, originally helped develop what iconic piece of Intel hardware?

- A] Intel's 80486 processor
- B] Intel's 8051 microcontroller
- C] The original Pentium
- D] The 8-bit 8008

18 Which classic piece of hardware was reborn this year, with AMD-powered internals?

- A] The Commodore 64
- B] The Sega Master System
- C] The Radio Shack TRS-80
- D] The Atari VCS

19 Nvidia has coined yet another acronym, CMP, what does it mean?

- A] Constant Modulated Power
- B] Cryptocurrency Mining Processor
- C] Color Modulation Process
- D] Cryptography Mining Protection

20 Which of these didn't Intel kill off this year?

- A] Consumer Optane memory
- B] The Itanium processor
- C] The Performance Tuning Protection Plan
- D] The fiber optic department

21 Have a guess at the top three tech companies in the world by market capitalization?

- A] Microsoft, Apple and Meta (Facebook)
- B] Apple, Amazon, and Alphabet (Google)
- C] Microsoft, Apple, and Alphabet (Google)
- D] Microsoft, Apple, and Tesla

More Hardware

22 What on a 7nm process IC is actually 7nm?

- A] The gate length
- B] Half the gate pitch
- C] The average track width
- D] Nothing specifically

23 We do like USB ports, don't we? What's the most you could have on an Intel Z690 motherboard?


- A] 10
- B] 12
- C] 14
- D] 16, and that's probably enough guys

24 Microsoft's HoloLens found a big new customer this year, who was it?

- A] The US Army, which wants headsets for the infantry
- B] The US Air Force to use in preliminary pilot training
- C] The CIA for 'unspecified' reasons
- D] The Federal Government to equip employees working from home

25 This year, the world's most powerful computer title was overthrown, by what?

- A] A recently revealed Chinese quantum computer called Zuchongzhi 2.1
- B] Perlmutter, which uses Epyc CPUs, and Ampere A100 GPUs
- C] Summit, an IBM machine with Power9 CPUs, and Tesla V100 GPUs
- D] Fugaku, a Japanese research machine with 158,976 processors in all



Soft Stuff

26 Google measures its Tensor AI chip's performance using a custom metric what is it called?

- A] Brain Floating Point Format
- B] Google Brain Floating Point
- C] Human Brain Equivalent
- D] Floating Brain Points

27 What's so special about Pragmatic's new version of the Cortex-M it is developing?

- A] It's printed on plastic
- B] It uses a ceramic substrate
- C] It uses a 2nm process
- D] It's going to be implanted into humans

28 Nvidia produced the GA106-302 variant to replace the GA106-300 GPU, why?

- A] To cure overheating with specific games that can cause permanent damage
- B] To limit the hash rate if the card was used for Ethereum mining
- C] To boost the hash rate if it's used for Ethereum mining
- D] To improve yields and help relieve the chip shortages

29 Just before Win11's launch, a survey asked users if they knew it was coming, how many did?

- A] About 62 percent of people said they knew about it, although most weren't ready to upgrade
- B] About half
- C] Only 38 percent of Windows users knew a new version was due
- D] Just 20 percent of people had a clue

30 Which OS do the world's top 500 supercomputers employ?

- A] A messy mix of commercial and custom software, no two are exactly alike
- B] Mostly versions of Unix, mixed with custom software and some Linux
- C] They all run tweaked versions of Linux, every single one of them
- D] Supercomputers require specially-written software specific to the hardware fitted and job in hand

31 What happens when you press the Windows key and the period key?

- A] The Windows language Settings appear, so you can switch easily
- B] A box with selectable emojis, kaomoji and unicode symbols pops up
- C] Playback on all video or sound files stops, press again to start them
- D] If you have a microphone attached, it will start recording but nothing appears on screen (spooky)

32 G Suite and Microsoft Office dominate in the US, but which is more popular?

- A] Big business still prefers Office, which has nearly 73 percent of the market
- B] It's fairly close, although Office has a moderate lead at 57 percent
- C] It's close, but G Suite just pips Office with 53 percent
- D] G Suite has grown rapidly since the growth of cloud servers, it now has 64 percent of the market



Soft Stuff *cont'*

33 What do the latest Microsoft Windows updates insist you remove?

- A] All versions of iTunes prior to 10.7
- B] Clippy, the Office assistant, if for some reason you still have him
- C] All versions of Netscape
- D] All versions of Adobe Flash

34 So how do you think Bing is doing these days?

- A] Surprisingly well. In September, it was the 11th most popular website in the country, beating Twitter
- B] Okay, it just makes the top 20, roughly level with Twitch
- C] Not terribly well, it's barely in the top 30 these days
- D] Oh, it's doing much worse than that

35 What's the most popular PC application, according to data from Google Trends?

- A] Currently, I think its Zoom Meetings
- B] Something utilitarian I guess, how about Microsoft Word?
- C] Spreadsheets are everywhere, so make that Excel
- D] Freemake Downloader is extremely popular

Bing zoom



Internet Inquisitions

36 Which of these smart devices can you currently buy?

- A] A smart rolling pin with an app to select dough thickness
- B] A smart kettle, for remote boiling using an app
- C] A smart candlestick that texts you when the candle burns too low
- D] None of these are real, yet

37 How many fake accounts does Facebook remove every quarter?

- A] Quite a few, about 5 million
- B] Lots, at least 50 million
- C] It might be as high as 500 million
- D] It can't be over a billion, can it?

38 Big data means big data breaches, how much did LinkedIn lose this year?

- A] Only a few, it lost 10,000 data records this year
- B] A number of breaches totalling about three million records
- C] A massive attack this year stole data from 92 percent of its users, over 700 million records
- D] LinkedIn has managed to remain a reasonably secure service

39 Smart phones are important, and addictive, how long does the average American spend using them?

- A] At least an hour a day, that's over a week per year
- B] It's probably nearer twice that
- C] It is over three hours a day, nearly two months out of your year
- D] Oh, it's much worse than that

40 Kids love TikTok, and it has made its own star performers, how many followers do the top six have?

- A] Up to 10 million each
- B] About 30 million each
- C] At least 50 million each
- D] Good grief, is it more than that?

41 Bitcoin mining continues to be popular, but how much electricity does the process consume?

- A] Compared to the world's total electricity consumption it's hardly anything to worry about
- B] It's a factor, mining eats through enough power to run a small city
- C] The mining uses about 0.5 percent of world power, that's roughly equal to a small country
- D] It's highly inefficient, and it is currently using over 1.3 percent of all the world's electricity

42 Which member of NATO did Google threaten with, and then briefly action, the withdrawal of its services this year?

- A] Bulgaria
- B] Norway
- C] Australia
- D] Belgium



Weird Science



43 To test the theory that tardigrades could survive space travel, scientists did what?

- A] Glued some on the outside of one of Musk's SpaceX rockets to see what would happen
- B] Built a miniature satellite to take them into orbit, currently awaiting a launch opportunity
- C] Dropped millions from weather balloons to see if they could survive the fall
- D] Fired them from a gun to test impact resistance

44 Google's quantum computer has been used to demonstrate the possibility of what phenomenon?

- A] Time crystals
- B] Time travel
- C] Quantum wormholes
- D] Invisible crystals

45 This year Google ended project Loon, what was it trying to do?

- A] Build an AI-powered robot drummer
- B] Map the surface of the moon
- C] Track common loon duck migration
- D] Use balloons for internet access

46 How much was paid for an NFT of the Nyan cat (basically, an animated gif)?

- A] \$56,100
- B] \$561,000
- C] \$888,888
- D] It wasn't more than that, was it?

47 Captain Kirk made it into space, how many people have managed this, as defined by the FIA?

- A] 392
- B] 556
- C] 675
- D] 1,243

48 One of these isn't a real thing, but which is it?

- A] The Microsoft Teletubbies plush toy
- B] The Intel Ottoman, a PC built into a footstool
- C] The iSmell, does what you think it does
- D] The Samsung Goyang, a phone inside a plush kitten toy

49 How big is the world's smallest 'submarine'?

- A] It has 244 atoms, so it's absolutely minuscule
- B] It's made from between 2,000 and 2,200 atoms, so it's very small
- C] Less than a millimeter, so it's about the size of a grain of salt
- D] A few millimeters, so it's about the same size as a grain of rice

GEEK POINT AVERAGE

Did you resist the temptation to check your answers?

Good. Now let's see how you measure up...

0-11 CORRECT: SUNNYDALE HIGH DROPOUT

There is something wrong here, you've managed to score below the average for random guesses. Look, even your cat could do better

12-13 CORRECT: MOS EISLEY COMMUNITY COLLEGE

This is the statistical average for blind guesses, so either you did just that or you have absolutely no clue at all

14-25 CORRECT: UNSEEN UNIVERSITY UNDERGRADUATE

So you do know a thing or two, keep reading *Maximum PC*, you have real potential

26-37 CORRECT: CAL STATE HOGWARTS

Well, you certainly have been paying attention! Our hearty congratulations, you are a geek

38-48 CORRECT: VULCAN INSTITUTE OF TECHNOLOGY

If you did this without help then we salute you; you've proved yourself worthy. Permission to feel very smug about yourself

49 CORRECT: YOU CHEATED, DIDN'T YOU?

We are pretty sure you did, although we have no proof

ANSWERS

Processor Posers 1]C, 2]A, 3]A, 4]B, 5]D, 6]D, 7]B, 8]A, 9]B, 10]A, 11]C, 12]B, 13]A, 14]D, 15]B, 16]A, 17]A, 18]D, 19]B, 20]D, 21]C, 22]D, 23]B, 24]A, 25]D, 26]A, 27]A, 28]B, 29]C, 30]C, 31]B, 32]B, 33]D, 34]A, 35]C, 36]B, 37]D, 38]C, 39]D, 40]C, 41]C, 42]C, 43]D, 44]A, 45]D, 46]B, 47]B, 48]D, 49]A.

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.
- Simplify your storage with one of our stylish binders, perfect for displaying on your shelf!



GIFT IDEA!
Gifts that
keep on
giving

VISIT OUR ONLINE STORE AT
www.magazinesdirect.com

✓ Convenient home delivery ✓ Best price guarantee ✓ Cancel any time ✓ 1.4m subscriptions sold

SUBSCRIBE TO **MAXIMUM PC**



GREAT REASONS TO SUBSCRIBE

- Save over 68% when you subscribe.
- Free *PC Hardware* digital bookazine.
- Never miss an issue.
- Get every issue delivered direct to your door.

FREE BOOKAZINE
WHEN YOU BECOME A
MAXIMUM PC SUBSCRIBER



TO ORDER GO TO
www.magazinesdirect.com/mpc/B2AMAG

TERMS AND CONDITIONS Digital bookazine only available with a paid one- or two-year subscription. Bookazine will be emailed 3-5 days after payment is processed. Savings based on the US newsstand cover price of \$9.99 per issue. Pricing valid on US orders, visit us at magazinesdirect.com for international pricing and other subscription options. Your first issue will arrive 4-6 weeks after your payment is processed, please allow additional time for overseas deliveries. Offer expires: 1st June 2022

LINUX WINDOWS

Old rivalries have been forgotten and ancient boundaries blurred. *Jonni Bidwell* explains the Redmond-Penguin harmony

BACK IN 2015, Microsoft decided that it loved Linux so much it built a whole Windows Subsystem for Linux (WSL), which enabled Linux programs and development stacks to run natively. The official statement even had a heart emoji.

Heart or not, some were skeptical of Microsoft's intentions, with memories of the "Embrace, Extend, Extinguish" mantra and the 1995 Halloween papers still fresh in their minds. But it seems as though Microsoft wants to accommodate Linux users (well, developers mostly), rather than force a mass defection.

A successor, WSL 2.0, was announced in 2019, built around a real Linux Kernel, rather than a Wine-like translation layer. It brought faster performance, swifter filesystems, and increased application

compatibility. In April 2021, a new feature was announced, WSLg, which enables graphical tools to run seamlessly on WSL. No need to shoehorn an X server on Windows, nor to redirect PulseAudio—it even works with Wayland. Not only can you run Bash on Ubuntu on Windows, but also Blender, GIMP, and Krita.

We'll see how to set this up on Windows, how it's a great way to learn Linux, and how to do some weird and wonderful stuff with it. Finally, we'll look at the efforts Linux has made to bridge the divide with its proprietary counterpart. So, if you have ever broken filesystems, networking, or bootloaders, everything should just work now with minor configuration changes.



PLATFORMS, TRUSTED PLATFORMS AND WINDOWS 11

Windows hardware requirements have always promoted confusion, and this time is no different...

WHEN WINDOWS 8 was announced, there was a concern that Microsoft would use Secure Boot, an optional feature of the then-new UEFI (Universal Extensible Firmware Interface), to restrict the installation of other operating systems. Secure Boot only allows booting EFI images that have been signed by a key enrolled in the UEFI. Since almost all hardware ships with a public signing key from Microsoft, it's easy to see where this concern came from. However, to help Linux distros (or any software that needed its own bootloader) deal with Secure Boot, Microsoft used its magic key to sign a small program (they probably wouldn't sign something big and complicated) called Shim.

Shim is a first-stage bootloader used by many distros (including Debian and Ubuntu) to launch their own (second stage) bootloaders. Secure Boot works by each stage checking the signature of the following one before executing it, establishing a root of trust back to the original signing key. So distributions can embed their own keys into Shim, and have this distro-specific Shim package signed by Microsoft.

Shim can then check the GRUB EFI image, signed by the distro key, and boot can proceed. It's possible to sign kernel images too, as well as subsequently loaded modules and firmware, so if you trust cryptography and the Microsoft Signing Authority, then Secure Boot makes it really hard for malware to be loaded anywhere in the early boot process.

A QUESTION OF TRUST

Of course, not everyone's going to trust Microsoft, or indeed Secure Boot. And that's okay, because in 2013 when Microsoft stipulated that "Windows 8 Ready" hardware should ship with Secure Boot enabled, it also made the provision that it should be possible to disable it, and also that it should be possible for users to install their own MOKs (Machine Owner Keys) and use them to sign whatever bootloader they wanted, so that ultimate trust, but also ultimate responsibility, lay with the user.

"Windows 10 Ready" hardware was subject to similar conditions, except the requirement that Secure Boot is mutable was reduced to a suggestion. Still, we've never come across a single Windows 10 machine where Secure Boot could not be disabled. And if you build your own systems, then you have nothing to worry about.

The "Ready" conditions are for OEMs that ship systems with the OS already installed. With Windows 11 upon us, it's no longer Secure Boot clauses that are being scrutinized, but those relating to TPM chips.



Trusted Platform Module (TPM) chips are tiny processors that have a small privileged memory store, which applications can use to store keys, authentication data, or anything that they don't want other applications nosing at or tampering with. Besides the minimum hardware specifications, for a PC to officially support Windows 11, it must also support TPM 2.0. At the moment, it's still possible to install Windows 11 via the official ISO if your device doesn't meet the TPM (or other) requirements, but Microsoft says such installations are unsupported and may not be privy to security updates down the line. Indeed, we have seen builds of Windows 11 running on a Pi 400, as well as a Nokia Lumia 950XL, one of the last devices to run Windows Phone.

Tiled start menus and the Windows desktop may bring you no joy, but look – Ubuntu can cheer you up

GIVE ME ALL YOUR TPM

TPM chips have been around for most of the past decade. TPM 2.0 was introduced in 2014, and most motherboards from 2016 onwards include one. TPM can be implemented in firmware too (so-called fTPM), at a slight cost to security, since some new Spectre/Meltdown-type attack could, in "theory", be leveraged against it.

Still, fTPM is good enough for Windows 11's requirements. You might need to enable TPM via the UEFI (classic BIOS is also not supported), where it goes by so many names that Microsoft made a friendly help page (see <https://bit.ly/lxf282-mshelp-tpm2>)

TPM is fully supported on Linux and can be used to secure SSH keys (see <http://blog.habets.se/2013/11/TPM-chip-protecting-SSH-keys---properly/>), unlock LUKS encrypted volumes (via `systemd-cryptenroll` or `Clevis`) or even make Secure Boot even securer (<https://threat.tevora.com/secure-boot-tpm-2/>). There are separate software stacks for TPM 1.2 (TSS aka `TrouSer5`) and TPM 2.0 (`tpm2-tools`), and there's a nice summary of both on the Arch Wiki page at https://wiki.archlinux.org/title/Trusted_Platform_Module.



LINUX IN WINDOWS

Never mind Windows 11, the latest version of the Windows Subsystem for Linux is where it's at



ARE YOU PART OF the slightly sinister-sounding Windows Insiders Program? Have you installed an even more sinister-sounding Preview Build of Windows 10 (at least build 20262)? Then it's easy to run Linux as part of WSL 2.0: just open an administrator-privilege Windows command prompt and run `wsl --install`. This also works on Preview Builds of Windows 11. A few clicks and pops and a restart later and you'll be in business.

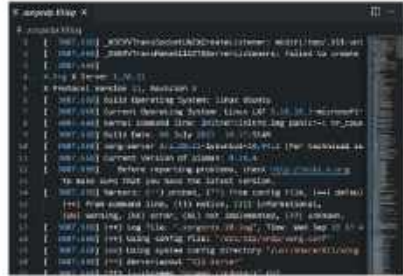
Actually, the GUI App Support download is pretty large, so now might be a good time to make a cup of tea. On your return, WSL, Microsoft's Virtual Machine Platform, their custom Linux Kernel and Ubuntu should all have been downloaded. Other distros are available and you can specify, for example, `wsl --install -d Fedora` to install Fedora instead, or additionally if you already have WSL Ubuntu installed. Users of non-Preview Builds can either join the Insiders Program and upgrade to one, or follow the manual installation steps below.

Start by firing up Administrator-powered PowerShell, then acknowledging that the command line can get ugly on Windows too, enter the incantation:

```
> dism.exe /online /enable-feature /
featurename:Microsoft-Windows-Subsystem-Linux /
all /norestart
```

This will install WSL 1, at which point you might wish to reboot and play with that. Or you could continue to get the latest incarnation. If you're not running Build 18362 (check by running `winver.exe`) or higher (or 19041 or higher form ARM64) you'll need to persuade the Windows Update Assistant to get you there. To enable WSL 2 we first must enable the Virtual Machine Platform, which requires this line to be entered at an

It's time to upgrade Windows 11, by adding a Linux kernel and some Microsoft magic



Now you can pore over logs and try and see what went wrong with RDP

Administrator PowerShell:

```
> dism.exe /online /enable-feature /featurename:VirtualMachinePlatform /all /norestart
```

One restart later and you'll be running the subsystem sequel. To get Microsoft's latest frankenkernel, download and run the package at https://wslstorestorage.blob.core.windows.net/wslblob/wsl_update_x64.msi. This is a separate package for ARM devices, so change `x64` for `arm64` in the previous link if that applies to you. Before we get to installing Linux we need to set WSL 2 as the default version from PowerShell, easy enough:

```
> wsl --set-default-version 2
```

Then get over to the Microsoft Store (<https://aka.ms/wslstore>) and install a distribution. The recognizable versions are all free, but if you're feeling flush you might want to try the Microsoft Research's WSL-tailored Penguin distro, currently on sale for about half of its \$20 price tag.


CREATE A STANDALONE ACCOUNT

Whatever you choose, and however you got there, the first step post-install is to follow the prompts and set up a username and password. This is entirely unrelated to any other accounts you have set up on Windows, and also separate from WSL accounts set up by other Windows users on the same machine. In other words, WSL machines are effectively single-user affairs, since they're the concern of a single Windows user. Then just as if it were a real Linux machine, you should do a system update:

```
$ sudo apt update
```

```
$ sudo apt upgrade
```

Where you go from here is up to you. You could just install Vim and the build-essentials package and get hacking on your latest C project. Or you might want to install the Nvidia Container Toolkit, which sets up everything you need to run GPU accelerated Docker



images, so you can use CUDA, PyTorch, and TensorFlow to machine-learn your way to freedom.

Or you might want to try something else. One of the best things about WSL 2.0 is that it supports GUI tools without any extra configuration. It even has its own Wayland compositor. In the previous version, graphical tools could be used but needed an X server (GWSL, VcXsrv, or X410) to be running on the Windows side.

Our favorite GUI tool is Microsoft Edge, and our favorite type of humor is sarcasm. So if you wanted to run the Linux edition of Microsoft's browser in a Microsoft-flavored version of Linux, you could go right ahead and do:

```
$ sudo apt install software-properties-common apt-transport-https wget
$ wget -q https://packages.microsoft.com/keys/microsoft.asc -O- | sudo apt-key add -
$ sudo add-apt-repository "deb [arch=amd64] https://packages.microsoft.com/repos/edge stable main"
$ sudo apt install microsoft-edge-dev
```

Now Edge should appear in the start menu, or you can start it with a simple `microsoft-edge`.

In this way, you could also install such Redmondian wonders as Microsoft Teams or Skype. Or perhaps something that isn't so Windows-centric. Popular open-source programs are just that, so the likes of GIMP, Blender, VLC et al all have official Windows equivalents. If for some reason these don't work well, then now you have options. The Transmission BitTorrent client is perhaps an exception—it has a Windows port, but it's described as an early preview. You may have your own Windows torrent client preferences (uTorrent is popular), but if not you can enjoy the solid reliability of the Linux version. It's just a matter of:

```
$ sudo apt install transmission
```

Visual Studio Code (or the VSCodium fork) is becoming increasingly popular on Linux. And it's fast becoming the code editor of choice for Raspberry Pi users. We can install the Linux version in much the same way as we installed Edge above, but we can also do something different.

Recent VS Code versions support the WSL-remote extension, so you can install the Windows version (from <https://code.visualstudio.com>) and then have easy access to your WSL files.



For this to work ensure "Add to PATH" is checked in the Select Additional Tasks dialog at the start of the installation. Then from the Extensions marketplace, search for Remote Development and install the extension pack (which enables remoting to WSL, Containers, or over SSH). According to the docs at <https://docs.microsoft.com/en-us/windows/wsl/tutorials/wsl-vscode> you should now be able to start VS code from within WSL by opening the Ubuntu command line and running `code`. This didn't work for us, but we were able to use F1 to bring up the command palette in VS Code and search for `Remote-WSL`, which reveals some of its many functions. These include being able to open a folder in WSL, complete with its own Terminal.

If you run the `top` command within a freshly started WSL 2.0 instance, you'll see there's not much running—a couple of init processes, Bash and Top itself, and not a hint of bloat. This is a slight subterfuge because a lot of the magic happens in an invisible service virtual machine. This runs a containerized distro based on Microsoft's CBL Mariner Linux (developed for its Azure cloud) which handles Wayland, X, PulseAudio and beams it back into Windows via FreeRDP. Read more about WSLg, at <https://github.com/microsoft/wslg>.

Having a whole Marketplace inside a code editor isn't everyone's cup of tea, but you'll find WSL extension fortifying and refreshing

FOREIGN FILESYSTEMS

From the Linux side, your Windows C: drive and everything in it is available at `/mnt/c` and from the Windows side, you can see your Linux filesystem(s) from Explorer by following the Linux shortcut (it has a penguin icon) below OneDrive, This PC, Network, and whatever else. While handy, you shouldn't let Linux programs do I/O heavy operations on `/mnt/c`. Why not? Read on...

If you run the `mount` command, you'll see that `/mnt/c` uses some sort of exotic Plan 9 (9p) filesystem to mount a drvfs volume. You'll also see that the

root Linux filesystem lives on its own block device (`/dev/sdc` or such) and that there are mount points for WSLg (the helper distro that enables tools to run seamlessly with graphics and audio).

In WSL 1, drvfs is used directly, which means that performance between Windows and Linux file systems was a little more performant. WSL 2's use of a 9p server to wrangle these transfers slows things down but also enables more flexibility and security. For this reason, WSL 1 will outperform WSL 2 at 'across the boundary' transfers.

So if you need to run, say, a database in WSL but, for some reason, its files must be visible in Windows, then you should downgrade to WSL 1. It's better to try and keep heavy transfers within the WSL bubble. It's even possible to add additional drives to your machine and format them as Ext4 (or your favorite Linux filesystem) from within WSL.

You won't be able to see them from the Windows side (at least not without additional tooling), but transferring large files or doing lots of random I/O should be much swifter.



YEAR OF THE LINUX DESKTOP (ON WINDOWS)

Running graphical applications is impressive, so the next step is to try and run a complete desktop. What are you waiting for?

X410 is a well-regarded and occasionally discounted X Server for Windows 10. We need an X411, though

GETTING AN ACTUAL Linux desktop running on WSL 2.0 on Windows 11 proved a little tricky. In fairness, WSL was built to run applications, not to act like a conventional virtual machine. It's not really surprising that GNOME and KDE won't work. They rely on all sorts of daemons and services running in the background holding all the components together. Brave souls have managed to get the Xfce desktop working nicely on Windows 10 with Ubuntu 18.04, but this didn't work for us on Windows 11 with Ubuntu 20.04. If you're running Windows 10, then you'd be well advised to try out running Xfce via the X410 server. It should just be a matter of adding:

```
$ sudo apt install xfce4
```

Then setting the `DISPLAY` environment variable correctly. A major change in WSL 2 is that `127.0.0.1` in WSL doesn't refer to the real loopback interface in

Windows (where the X410 server would be listening). So we need to figure out the WSL address and update `DISPLAY` accordingly. As per the instructions at <https://x410.dev/cookbook/wsl/using-x410-with-wsl2> this can be gleaned from `/etc/resolv.conf` on WSL, and then all that's left is to tell X410 to allow "Public" connections. You will probably want to use Windows Defender Firewall to restrict this (if your machine isn't otherwise firewalled). Then run `xfce4-session` within WSL.

TAKING A DIFFERENT APPROACH

While not technically in the remit of WSL, it's still nice to have a Linux desktop on Windows 11—if for nothing else than to show how the new-fangled desktop measures up to something tried and tested. Since using a native X server on Windows didn't work, we figured we'd try a different approach: using Windows' own Remote Desktop Protocol (RDP) to connect to a desktop running within WSL2. To do this, we installed Xfce in Ubuntu using the command above [note that in some situations you may need to install `xfce-terminal` explicitly as well], and then installed and started Xrdp:

```
$ sudo apt install xrdp
```

```
$ sudo /etc/init.d/xrdp start
```

Note the old Sys-V style init script. WSL doesn't use Systemd, which will surely be good news to some people, but is also part of the reason you can't use GNOME or it right now. Anyway, we're getting sidetracked. Start Remote Desktop Connection in Windows and connect to `localhost:3389`. You'll be warned that the identity of the remote computer cannot be verified, but since the remote computer here is, in fact, local this is nothing to



HARDWARE-ACCELERATED WSL

WSL supports running with a virtual GPU so that Linux programs can take advantage of high-powered graphics cards and run compute jobs (not games at this stage) at near-native speeds. There's also a DirectML backend for TensorFlow (see <https://github.com/microsoft/DirectML>), enabling speedy machine learning experiments on Direct3D 12-supported hardware.

Further to this, in spring 2021 the `d3d12` driver became an official part of the open-source Mesa 21.0 stack. To make use of this, and speed up your Blender renders or whatever, you might


need to update your Windows GPU driver to one supporting WDDM 3.0.

Nvidia's WSL CUDA drivers are considered public preview and should be delivered to Windows Insiders through the normal Windows Update channel. Otherwise, they are available from <https://developer.nvidia.com/cuda/wsl>.

Owners of AMD graphics hardware can get the preview drivers at www.amd.com/en/support/kb/release-notes/rn-rad-win-wsl-support. Intel's can be found at <https://bit.ly/lxf282-intel-drivers>.

Only Ubuntu LTSes are available on WSL, they miss out on all these developments. So you could install Arch or some other distro. Or you could avail yourself of Canonical's Ubuntu on Windows Community Preview. Read more and follow the download links at <https://ubuntu.com/blog/announcing-ubuntu-on-windows-community-preview-wsl-2>.





Eventually our efforts paid off and we were greeted with the minimal, yet glorious Xfce desktop

worry about. A login screen should open, from which you should select the Xorg session and use your WSL credentials to log in. At first, this didn't work for us, and we got nervous because time was getting short.

We encountered a mildly encouraging blank screen, which then vanished, dashing our hopes. Then we remembered that X sessions are complicated things, so we had little expectation that the following hack would actually work. But it did! And hopefully, it does for you, too. Make a trivial session file in your home directory (on WSL) as follows:

```
$ echo xfce4-session > ~/.xsession
```

And now try and log in again. If it doesn't work, you can always log out of WSL, shut it down and restart it by running:

```
$ logout
> wsl --shutdown
> ubuntu
```

Our session seemed to die whenever the Windows screensaver activated, necessitating this step on a number of occasions. Once you get it working, you'll notice that the image quality isn't quite pixel-perfect, but this can be tweaked by giving the server more bandwidth. Exit Xrdp and edit the config file with:

```
$ sudo nano /etc/xrdp/xrdp.ini
Next, find the max_bpp value,
change it from 32 to 128. Below this
add the line
xsserverbpp=128
```

Now restart Xrdp with:

```
$ sudo /etc/init.d/xrdp restart
```

and you should have a slightly sharper, slightly more responsive desktop.

LEARNING TOOL

If you're something of a Windows expert, but don't have much familiarity with Linux, then having this remote desktop on hand will be a great learning tool. It's a hip alternative to a virtual machine. And just like a virtual machine, you can export the installation and run it on another machine, or on the same one if you want a handy backup before you try some adventurous command line juju. To export a WSL image as a tarball run the following from a Windows Command Prompt:

```
> wsl.exe --export Ubuntu c:\ubuntu.tar
replacing Ubuntu if you used another distro. You
can keep track of which WSL distros you have installed
with
```

```
> wsl.exe --list --all
```

Then, if you have cause to re-import a tarball you can give the distro a different name:

```
> wsl.exe --import UbuntuTweak c:\mywsl c:\
ubuntu.tar
```

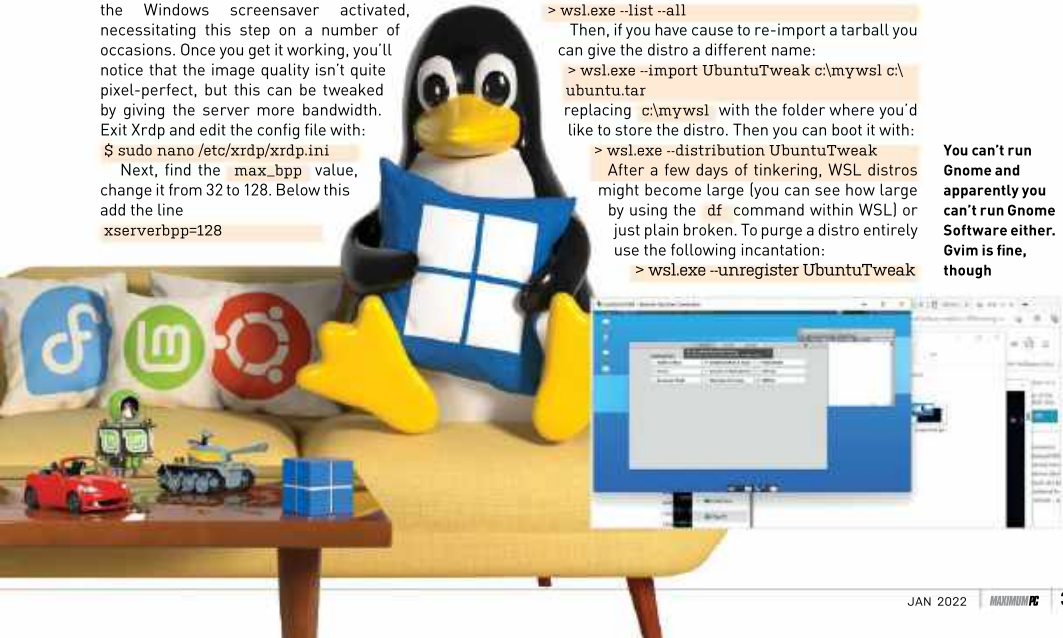
replacing c:\mywsl with the folder where you'd like to store the distro. Then you can boot it with:

```
> wsl.exe --distribution UbuntuTweak
```

After a few days of tinkering, WSL distros might become large (you can see how large by using the `df` command within WSL) or just plain broken. To purge a distro entirely use the following incantation:

```
> wsl.exe --unregister UbuntuTweak
```

You can't run Gnome and apparently you can't run Gnome Software either. Gvim is fine, though



DUELING OSes

Linux has always tried to co-operate with Windows oddities and, despite common gripes, it does a pretty good job

OWNERS OF A MODERN UEFI system should be able to revert any unwanted changes caused by installing a new operating system, as long as you can get into the UEFI settings. The UEFI boots an image from the EFI partition or live media, and userspace tools (both on Windows and through efimg on Linux) can change the boot order and, in some cases, break things. It's generally expected, but not necessarily desired, that a newly installed OS will put itself at the top of the boot ordering. Often this results in Windows booting immediately and the only way to get to the UEFI settings to revert this is to hold the Shift key while shutting down from the start menu. An option to reboot to settings should appear.

Our experience was, mercifully pain-free though. We upgraded Windows 10 on a 10th generation Dell XPS set up to boot Ubuntu, and lo and behold the default boot choice (GRUB) was left intact. Not only did Ubuntu boot correctly, but the old GRUB entry for chain loading Windows also worked. We think we got lucky here, but even if we didn't, booting Windows would just be a matter of entering the UEFI settings or summoning a boot menu at startup. If UEFI didn't give us these options, then we can reboot to the firmware setup with `$ systemctl reboot --firmware-setup`

For older BIOS machines the situation is different. Each hard drive can have its own Master Boot Record (MBR), and any OSes installed can put their bootloaders here. Well, the first part of their bootloaders, the MBR is only about 440 bytes long, and modern bootloaders like GRUB are complicated, so the MBR contains a skeleton program that tells the machine where on the hard drive the rest of the bootloader can be found.

Grub version 1 also had a Stage 1.5, which filled about 30KB directly after the MBR and contained the filesystem driver for the partition containing Stage 2. Grub v2 doesn't need Stage 1.5, as Stage 1 points directly to the logical block address containing Stage 2 (without worrying about filesystems), and the Stage 2 will load the main config file `/boot/grub/grub.cfg` as well.

Modern UEFI makes it difficult for Windows to prevent Linux from booting



We have long recommended keeping separate OSes on separate drives, and one of the main reasons for this is to mitigate against one OS's bootloader treading on the toes of another. This

is as much a problem between different Linux distros as it is between Linux and Windows. Linux distros are generally good enough to give you a boot menu, but entries for secondary distros won't respect any of their settings in `/etc/default/grub`, which may cause issues.

Also, if one of these distros is Arch Linux, it often doesn't get detected by the `grub-install` scripts. Installing the `lsb-release` package on the Arch side should solve this. Windows has no qualms about obliterating any other bootloaders on its drive. It's easy enough to get GRUB back when this happens, but better to avoid it happening by giving Windows its own drive.



ROGUE WINDOWS UPDATES

But this might not be enough. We've read plenty of posts online saying that Windows updates have boldly gone where they're not supposed to go, and we've been around long enough to have seen this behavior firsthand. Planting recovery partitions on other drives, corrupting LVM and RAID volumes, and all kinds of other unorthodox behavior. Windows has nothing to gain from this, but living alongside another OS was never in its remit, so in a sense, this is to be expected.

Besides battling bootloaders and (maybe) dodgy `.docx` files, one of the major points of friction between Windows and other OSes is the NTFS filesystem. It's not just a concern for Linux—getting macOS or Android to write to these volumes was always tricky too, at least without proprietary tools. On Linux, we at least had the userspace NTFS-3G driver, but this was not the most performant. Mercifully, help is at hand in the form of a new kernel driver (see box, right).

Also, we shouldn't bash NTFS-3G. The official WSL documentation warns that the converse situation, viz. transfers from within WSL to without won't be blazing fast either. This is a slight regression from WSL 1, which thanks to not having a real Linux Kernel was able to streamline transfers across filesystems. Still, certain I/O-heavy workloads to Linux filesystems that were hobbled under WSL 1, should now do much better.

Besides NTFS (see below), Paragon Software has an agreement with Microsoft that enables it to market and license its own exFAT implementation[s]. This would have been something of a boon for Paragon, but

Microsoft effectively open-sourced exFAT in 2019. More precisely, it made public the filesystem's specifications and offered free licensing to any member of the Open Invention Network (OIN), the holders of a defensive patent pool to which Microsoft and others have contributed their IP. It left the actual work of writing a driver to the community, and this, thanks to Samsung, was added to Kernel 5.7.

Again, there was previously an exFAT kernel driver which first appeared in version 5.4. This was based on an old driver for one of Samsung's Android tablets, the code accidentally ended up on GitHub in 2013 and was later open-sourced by Samsung. That code ended up in the Kernel's Staging tree where it was improved, before being replaced by Samsung's newer code drop.

Network filesystems have been improved by Samsung too. Kernel 5.15 includes a new driver, KSMDB, for serving shares over the latest version of Windows' Server Message Block (SMB3) protocol. This should make for faster transfers in the immediate term, but long term, the goal is to have a leaner project that's easy to add new features to as the protocol evolves.

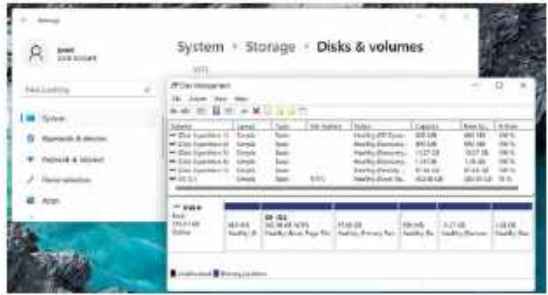
This is in contrast to the userspace Samba, a massive project that deals not just with serving files over SMB, but also client tools, authentication, Active Directory, and much more. So KSMDB is designed to work in harmony with Samba, rather than replace it.

High on KSMDB's list of things to do is to implement RDMA (Remote Direct Memory Access) support, also known as SMB Direct, which enables file transfers to skip protocol overheads, which is currently a bottleneck for both mobile and embedded systems.

Samba is vital for anyone who wants to use a Linux machine on a Windows network and for the most part it works out of the box. There are some edge cases though. For some reason, plenty of people have run into the issue of old 'LANman' password authentication no longer being supported, so out of the box, it's not possible to connect to SMB shares on Windows XP (or older) machines.

These machines also, by default, use an older version of the protocol, and to make Samba play along the generally accepted solution is to add the following lines to `/etc/samba/smb.conf`:

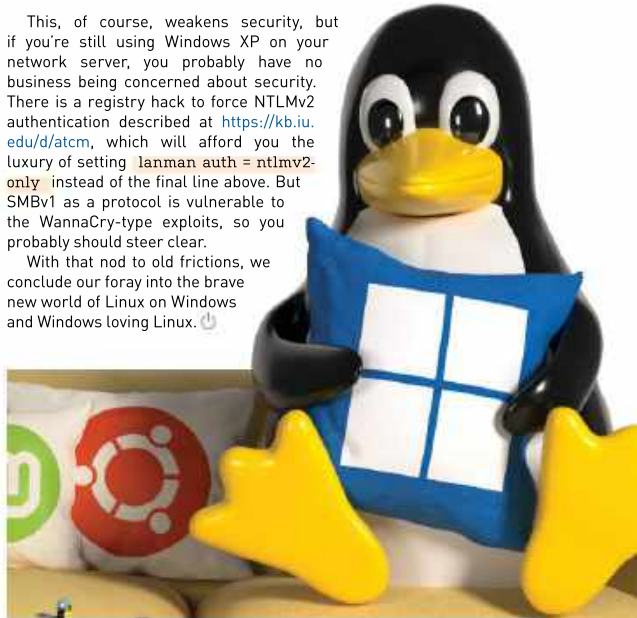
```
server min protocol = NT1
lanman auth = yes
ntlm auth = yes
```



Windows Disk Management tool begrudgingly accepts our Linux partition is there, but offers no info

This, of course, weakens security, but if you're still using Windows XP on your network server, you probably have no business being concerned about security. There is a registry hack to force NTLMv2 authentication described at <https://kb.iu.edu/d/atcm>, which will afford you the luxury of setting `lanman auth = ntlmv2-only` instead of the final line above. But SMBv1 as a protocol is vulnerable to the WannaCry-type exploits, so you probably should steer clear.

With that nod to old frictions, we conclude our foray into the brave new world of Linux on Windows and Windows loving Linux. 🐧



PARAGON NTFS3 DRIVER

After years of having to make do with the NTFS-3g FUSE (Filesystem in Userspace) driver to read and write to Windows NTFS volumes in Linux, a driver is finally coming to the kernel. Actually, there has been a kernel driver for a long time (since 2001 in fact), but it only ever provided read-only support and was overlooked by most people.

NTFS-3g was a valuable crutch for people who worked on dual-boot

systems but compared to native NTFS its performance was weak. On older systems, in particular, it would consume a lot of CPU cycles for the underwhelming rate at which it wrote bits.

The new kernel driver, dubbed NTFS3, comes from Paragon Software, a commercial provider of cross filesystem, partitioning, and network management tools. NTFS has been around since 1993 and is of much

less commercial interest than it once was. Microsoft has superseded it in new versions of Windows, so Paragon chose to open source it last year.

This was a rocky road: Paragon's first attempt to dump 27k lines of spaghetti code to the Kernel was met with a swift refusal. However, 22 tries and some community advice later, the code has finally landed in Kernel 5.15, which will be released by the time you read this.



STEP ONE in hardening the security of your home network lies in determining if your router remains fit for purpose. If it's over six or seven years old, now is a good time to consider upgrading to something more robust that supports the latest security protocols. The boxout on page 45 reveals what to look for in a new router—the Synology RT2600ac we recommend is one of two routers we've tested for this feature, the other being an older—and more consumer-oriented—Archer VR900 modem router. Alternatively, check out the box on DD-WRT on page 46 to see if it's possible to give your existing model a new lease of life.

Whether you've decided to splash out on a new router, or simply wish to perform a network security audit with your current model, step one is to log into your router's management tools, which is typically done through your web browser. All you need to do is type your router's IP address into the browser and hit Enter. If you're wondering what its IP address is, open Settings → Network & Internet in Windows 11 and click Properties. Your router's IP address will be listed next to IPv4 gateway.

You'll be prompted to log in, typically with a username and password, but sometimes just a password. If you've never done this before, then consult your router manual. You'll discover the default login details aren't exactly secure, and this is where you should perform your first tweak: change them to something infinitely more challenging to bypass.

Look for an admin section in your router's management utility. Archer users will need to navigate to Advanced → System Tools → Administration to change this to something more secure; Synology users would change this during the initial setup, or later on under 'Control Panel → User'. Synology's approach scores points because it adopts a multi-user model, so would-be hackers would need to guess your username as well as your password to gain access.

We recommend using a password manager's generator such as that offered by Bitwarden to create your new password—something random, containing a mixture of letters, numbers, and



Make sure you choose WPA2/WPA3 for Wi-Fi encryption for backwards compatibility

special characters, and at least 14 characters long to protect it from brute-force attack. Obviously, make sure this is saved in your password manager or written down somewhere safe.

While you're in the administration section, look for an option that allows you to administer your router remotely—in other words, from outside your local network over the internet. Disable this feature to prevent any drive-by hackers from attempting to log in.

If you think there will be times that you'll need this feature, look to strengthen your account further. In the case of Synology routers, you can apply two-step verification via a 2FA code to add an extra layer of security. Try to get into the habit of leaving the feature disabled, except when you need it (say when traveling), to minimize its exposure to the internet.

CHANGE YOUR NETWORK SUBNET

Now is also a good time to consider changing your network's default subnet. Most routers default to

192.168.0 or 192.168.1 for their subnet and assign themselves an obvious IP address, such as 192.168.1.1 at the same time. This makes IP addresses easier for hackers to guess, even after you've disabled the automatic assignment of IP addresses to new devices via DHCP (see further down). Why not pick something less obvious, for example, 192.168.42 as a subnet, and 192.168.42.187 for your router's IP address?

Be warned: if you've already assigned manual IP addresses on the router's current subnet to certain devices for whatever reason, they will lose network access as soon as the change is made, so switch them back to automatic/DHCP for now. If you're looking at your Windows 11 PC, for example, you can do this via 'Settings → Network & Internet'. Click your network adapter (WiFi or Ethernet) followed by Edit next to 'IP assignment' → set it back to 'Automatic (DHCP)'.

Once done, make the change: For example, Synology users should head over to 'Network Center → Local Network → General tab' where you'll see your router's current IP address under Local IP. Change this to the desired replacement, leaving the Subnet mask set to 255.255.255.0. Make sure you update the DHCP settings beneath this too.

After clicking Apply, you'll momentarily lose your connection, but when it returns, you should be able to log on to your router again via its new IP address to continue tightening security.

Disable WPS when it's not needed



SECURE YOUR WI-FI NETWORK

With your router's own administration portal locked down, it's time to turn your attention to your physical network. In other words, we're going to secure both its wireless and wired connections to prevent any hackers within the vicinity of your network from gaining access.

Let's start with your Wi-Fi network. Navigate to its Wireless settings to see what encryption you've applied. As a bare minimum, it should offer WPA2-PSK with support for AES (avoid the more vulnerable TKIP encryption). Make sure WPA2-PSK/AES is selected if this is the case.

Modern routers support the newer WPA3 standard, Synology added support for this via a firmware update, but not all your network devices will. WPA3 support is required at both hardware and OS level, so even if your computer's Wi-Fi adapter supports WPA3, you'll need to be running Windows 10 or 11, or Ubuntu 20.04 or later. Compatible Apple devices include iPhone 7, iPad 5, Watch 3, Apple TV 4K, or Mac from 2013 with 802.11ac support. Android devices need to be running Android 10 or later.

Don't forget all your other wireless devices to: your TV, consoles, Roku box, and so on. It's likely you won't be able to switch exclusively to WPA3, but in anticipation of this, your WPA3-compatible router should offer you a backward-compatibility setting that offers WPA3 to those devices supporting it, while falling back to WPA2-PSK/AES for those that don't. In the case of the Synology, configure this via Wi-Fi Connect → Wireless → Wi-Fi section: Select WPA2/WPA3-Personal from the dropdown menu.

CHANGE AND HIDE YOUR SSID

Your network 'advertises' itself to others using its SSID, and there are two ways in which you can strengthen this part of your network. First, change the SSID name—the default usually includes the router's model or manufacturer, which can aid a determined hacker in gaining entry to your system by revealing what hardware they are up against. You can do this from the Wireless section of your router's configuration utility.

UPGRADE YOUR ROUTER

If you've determined your router needs upgrading, now is the perfect time to put security at the heart of your purchasing decision. And that means looking beyond consumer routers for something more suited to a SOHO environment. These types of routers receive frequent firmware updates (as opposed to none in many cases) to close security loopholes and occasionally add new features, plus they feature useful extras such as built-in VPN servers, more granular parental controls, and firewalls.

While SOHO routers carry a premium, they don't have to cost the earth. Synology's RT2600ac (above) costs around \$200 from newegg.com, and while it's not the quickest, it's certainly a lot more



secure than any consumer router and is one of the most affordable routers we've seen supporting WPA3 encryption.

Like most SOHO routers, the RT2600ac has no built-in modem, so if you've previously been using a modem router, keep that as your modem going forward. Setup is relatively straightforward: consult instructions on placing your existing modem router into modem-only mode (if it doesn't exist, you'll need to examine ways of setting it up in bridge mode), then connect its WAN port to your router's WAN port to link them together.

SOHO routers are traditionally more complex to set up and use, which

is another reason why the RT2600ac is a good choice. It runs its own dedicated operating system, Synology Router Manager (SRM), which receives frequent updates and provides a relatively simple point-and-click interface, as well as allowing access via a mobile app.

SRM can also be extended via applications you download and install from the internet, including Security Advisor, an app that scans your router's security settings to reveal weak points and suggest improvements. We're pleased to report that after following the advice in this tutorial, we got a 'Good' security rating from our RT2600ac on both home and work baseline checks.

You can go further and prevent your network from advertising its presence by disabling the SSID broadcast signal (it's a simple 'Hide SSID' checkbox on the Archer VR900). This prevents your network from showing up when people nearby scan for Wi-Fi networks to connect to; instead, anyone wishing to connect to your network for the first time must manually type in the SSID when setting up the connection (thereafter, they will connect automatically as usual).



Make sure guest devices are isolated from your network

It's not a perfect solution as hackers can still easily sniff out the SSID using legitimate tools such as inSSIDer (www.metageek.com/products/inSSIDer/), but it does deter casual snoopers. We suspect the extra layer of security outweighs the inconvenience of manually entering your SSID each time you connect a new device to your network.

RESTRICT GUEST ACCESS

All modern routers offer a separate guest wireless network for visitors and other infrequent users, and it's something we strongly recommend you enable, again using the best encryption your router supports and with a strong password.

By default, the guest network should remain isolated from the rest of your network, so while visitors can access the internet, they can't use network resources such as shared folders or printers or peer into shared folders.

DD-WRT INSTALLATION

If your router is showing signs of its age, but you don't want to upgrade, then visit the DD-WRT webpage <https://dd-wrt.com/support/router-database/> and see if your router model is listed as being compatible with the free DD-WRT alternative firmware. If not, consider an affordable upgrade, such as the Archer C7 (\$44 for the AC1750 model) that can be flashed with DD-WRT.

DD-WRT adds useful functionality to consumer-grade routers, including additional security tools, such as a firewall, VLAN support, and a VPN server. It's also frequently updated, so your router stays up to date. Not all features are available on each device, so after confirming compatibility try Googling your model number and 'dd-wrt' to see how others have fared.

Although the main DD-WRT interface is straightforward to



DD-WRT works with a wide range of routers

navigate, some features require a degree of technical know-how, which is why you should take a deep dive into the DD-WRT wiki (go to <https://forum.dd-wrt.com/wiki/> and click Tutorials) to see what level of complexity certain features require—the OpenVPN server is one such example.

In addition, installing DD-WRT could also potentially 'brick' your router, so make sure you familiarize yourself with the procedure, take your time, and download a copy of your manufacturer's

firmware first in case something goes wrong. If you do 'brick' your router, you may be able to recover from it via a hard reset or by using TFTP to restore your original firmware.

You can access a TFTP client in Windows 11 via the 'Turn Windows features on or off' tool, and see https://forum.dd-wrt.com/wiki/index.php/Recover_from_a_Bad_Flash for details of recovery procedures if you need them. Remember, if you decide to install DD-WRT on your router, you do so entirely at your own risk.

Depending on your router, you can choose to allow guests to see each other on the guest network without giving them access to your main network, plus set a limit on the number of devices that can be connected simultaneously—consider reducing this to one or two, depending on who is visiting. Synology users can also set up a Guest Portal that can be used to set time limits on guests' usage of your network.

MAC ADDRESS FILTERING

Another common way to restrict what devices can access your wireless network is to set up MAC address filtering. Each device is identified using a unique MAC address—six two-digit hexadecimal codes such as '00:0a:95:9d:68:16'—that helps identify it when it's connected to your network.

A MAC address filter allows you to set up a list of allowed devices

Assign IP addresses manually and disable your DHCP server



based on their MAC address, which in theory, blocks any other device from accessing your network or the internet, even if they successfully connect to your Wi-Fi network. The settings are usually easy to find, Archer VR900 users, for example, should navigate to 'Advanced → Wireless → MAC Filtering'. Choose 'Allow wireless access only from the devices in the list below', then click Add to enable them one by one.

Your router should have already detected which devices are connected to your network, enabling you to identify them by name and assigned IP address to simplify the initial setup. If not, then use a free tool called Advanced IP Scanner (www.advanced-ip-scanner.com) to generate a complete list.

MAC filtering isn't bulletproof because it's possible for devices to spoof their MAC address, but for a hacker to gain access, they'd need to know which MAC addresses had been whitelisted and use a tool like Technitium MAC Address Changer (<https://technitium.com/tmac/>). It may not be a perfect solution but filtering the MAC address is still a stumbling block worth considering.

DISABLE DHCP

In a similar vein, you could also disable your router's DHCP server, which is responsible for allocating IP addresses to devices as they connect to your network. DHCP is almost like rolling out the welcome mat, and while convenient, you may feel it's a better idea to disable it altogether and rely on devices being manually configured with the three key things they need to connect: your router's IP address as the gateway, the subnet mask (255.255.255.0) and a unique IP address based on the same subnet as your router, so 192.168.x.y, where 'x' matches that of your router, but 'y' is a number between 0 and 255 not allocated to any other device.

If you go down this route, it's worth creating a list of all your connected devices (Advanced IP Scanner could help here) and then allocating IP addresses to each one before you apply them in turn. Taking your Windows 11 PC as an example, you'd navigate to 'Settings → Network & Internet' and then click Wi-Fi or Ethernet (for wired connections) followed by Edit next to 'IP assignment'.



Change it to Manual, then fill in the details required. If you want to override your router's DNS address, you can do so here too before clicking Save. We'd recommend you allocate manual IP addresses to all your devices before finally disabling the DHCP Server (Synology users will find this under 'Network Center' → Local Network → General tab', for example).

CLOSE OFF WPS

There's one gaping security hole still to close, and that's WPS (Wi-Fi Protected Setup). WPS is designed to simplify adding compatible devices to your network—you either set up a PIN or press a button that puts your router into WPS scanning mode. When you press the corresponding button on your device, it is automatically connected to your Wi-Fi network, no further questions asked.

While WPS is a useful timesaver, you shouldn't get into the habit of leaving it on as anyone with the PIN number or access to the WPS button on the router could use it to surreptitiously connect a device. Here, you could consider looking for the settings to disable both options when it's not in use. Look under 'Advanced' → 'Wireless' → 'WPS' on an Archer router, or open Wi-Fi Connect and navigate to 'Wireless' → 'WPS tab' on a Synology model, for example.

CHECK YOUR PORTS

Now your network is closed off to potential physical incursions, it's time to examine how to protect yourself from attacks originating online. Disabling remote administration and changing your

Create a whitelist of permitted wireless devices by MAC address

router's management passwords are good first steps, but let's go a little further than that.

To communicate with other network devices, as well as access services over the internet, your

router is obliged to open ports to allow traffic to move freely from one place to another. Ports make it easy for traffic to reach its correct destination and some are well-known, such as 80 for http web traffic, 20 and 21 for FTP, and 25 for opening an unencrypted SMTP connection to send email through.

There are theoretically 65,535 ports available, although the number available to services is closer to 49,000. Any activity involving a server, whether hosted locally or over the internet, will make use of these ports, and while some are official, such as 32,400 for Plex, others tend to be looser.

Before the advent of Universal Plug 'n' Play (UPnP), you'd be obliged to go into your router to set up these ports manually, specifying the port number, its protocol (TCP and/or UDP), and its destination on

DIAL HOME SECURELY

Have you ever been on the road and needed to access your home network in some way, such as copying files to or from a shared folder, or sending something to your home printer for when you return? If you use a VPN server, you can add this functionality securely without leaving your network open to attack.

VPN servers require existing hardware set up inside your home network.

This might be on your SOHO router (VPN Plus is part of Synology's SRM software), or you could add it to your server: PiVPN (<https://pivpn.io/>) works with a spare Raspberry Pi, for example, or you'll find support in higher-end NAS drives from the likes of QNAP and Synology.

You can even create a VPN server on your Windows PC, although you'd need to leave it switched on and connected while you were away to gain access—see [Helpdesk Geek \(https://helpdeskgeek.com/](https://helpdeskgeek.com/)



Synology offers a range of VPN server options

[windows-10/how-to-set-up-the-windows-10-built-in-vpn-service/](#)) for a handy guide.

When connecting through a third-party VPN service, we recommend using the L2TP/IPSec protocol for compatibility purposes, making sure you enable the 'Pre-shared key' option to force users to provide additional authentication before connecting.

You'll need to configure your laptop to access the VPN server when you're out and about. In Windows 11, navigate to 'Settings →

Network & internet → VPN' and click 'Add VPN'. Select 'Windows (built-in)' for VPN provider, then give your connection a descriptive name, and fill in the details required: server name or address, which is basically your public IP address (see www.whatsmyip.com) or domain or dynamic hostname if you've set one up, pre-shared key, plus username and password (see the Permission tab on VPN Plus Server on your Synology). And that's it, you'll be able to connect to your local network when away from home.



Parental controls allow you to protect vulnerable household members

your network (the IP address of the server or device traffic on that port should be redirected or routed to).

UPnP is a network protocol designed to allow compliant devices to set these rules for themselves, so when your device requests access to a specified port the rule is automatically set up. That's great for convenience, but not for security as once those ports are opened, they can be used for nefarious purposes as well as legitimate ones. Worse still, while UPnP only responds to requests for access from your local network, it assumes those requests are legitimate, so if malware gets on to your PC it can open ports without even a cursory check.

The simplest solution is to disable UPnP completely, but before doing so, you should review what ports have been opened, and assuming you recognize them, create manual rules to cover them. On your Archer, navigate to 'Advanced' → 'NAT Forwarding' → 'UPnP'. Make a note of what's on the UPnP Service List: description, external port, protocol, internal IP address, and internal port, which should help you identify what's there.

Switch to the 'Virtual Servers' section and click Add to add them one by one. The Service Name can be whatever you like to help identify it in the future, but be sure to copy the other information exactly as it's recorded. Once done, click OK and move on to the next. Once all your services have been set up, return to UPnP, and flick the switch off.

If you think you can't live without UPnP, then consider this



Close possible security holes by updating your router's firmware.

compromise: set a reminder to check it once a week for new services. Once vetted, repeat the process of converting them into virtual servers, then flick the UPnP off and on again to clear the list.

UPNP AND SYNOLGY

The process is slightly different with the Synology. Everything is conveniently located within a single screen (Network Center → Port Forwarding → Port Forwarding

tab), and you can combine multiple ports within a single port forwarding rule (if you do this, however, you must leave the 'Private port' field blank — it will automatically select all the ports you've defined under 'Public port').

Once you've created a port forwarding rule here, the corresponding UPnP rule will vanish, although it may reappear if another device on your network subsequently tries to access the same port. Unlike the Archer VR900, the Synology router has a built-in firewall that screens all traffic as it comes in from the internet.

However, by default, it has been configured to automatically create rules granting access to any ports that have been configured here, either manually or via UPnP. Click the Settings button to change this behavior. You might choose to disallow automatic rules for UPnP-created rules, for example, but you will have to remember to add these yourself whenever a new service attempts to open a port).

If you want to disable UPnP on the Synology router, switch to the Local Network section, scroll down to the DHCP Server section, and set the 'Enable UPnP' menu to Disabled.

SET PARENTAL CONTROLS

While they are no substitute for dedicated parental control software, your router will probably offer some basic parental control tools. The Archer VR900 is typical of most consumer routers in that its options are fairly limited: you can set time limits for specific devices based on their MAC address, and create simple content filters for websites to block by URL or keyword.



Manual port forwarding rules are more secure than UPnP

The Synology scores more highly in this regard: open Package Center and click Open next to the Safe Access parental control app. Rules are assigned to profiles, which can be individuals in your home (covering all their devices), all unassigned devices on your network, or the Guest Wi-Fi.

Once assigned, you can limit users' access to the internet to specific times of the day or set a maximum time limit per day. As with the Archer, there's a web filter that can block by category as well as by domain. It offers templates for children, guests, employees, and so on, which you can then further fine-tune. Finally, you can even define a page style that users will see when they fall foul of the filter.

Safe Access also provides Synology users with additional protection against phishing, malware, and potentially unwanted software via the Threat Intelligence database and Google Safe Browsing (API key required for the latter).

MORE TWEAKS

Updates are a key part of keeping us safe online, but it's a sad fact that most consumer routers rarely, if ever, receive firmware updates. Taking the Archer VR900 as an example, it comes in three hardware versions and, of those, V2 has received no firmware updates, while V1 and V3 have only received a handful. On the other hand, the Synology router's SRM operating system receives frequent updates and can be configured to automatically check for updates.

Whatever model of router you have, look for an option to check for updates within your router's configuration utility. If none are found, check the support section of your router's website in case there is one you can manually download and install instead.

Synology routers also offer a dedicated Security section under Network Center that is worth exploring. Note that the 'Enable DoS Protection' setting applies more to organizations than individuals, as you're unlikely to become a target unless running services directly from your home network. However, we recommend using the Auto Block tab to help frustrate brute-force attacks on the router's configuration utility.

PROTECT INDIVIDUAL DEVICES

Securing your router protects the gateway to your network, but it doesn't mean you can ignore the security of each computer or device you've connected to it. After all, your network is only as secure as its weakest element.

Your Windows PC should be protected by the built-in Windows Security, but what about your mobiles, tablets, and other connected devices? While iOS devices don't technically require security apps, make sure you're running the latest version of iOS and if you have installed third-party protection on your PC, look for a mobile app such as Bitdefender Mobile Security that can highlight potential vulnerabilities and provide some email and web protection.

Android devices require antivirus just like your PC, and it's good practice to protect Macs too. Perhaps now is the time to purchase a multi-device security



Identify connected devices and check security on each one

package like Bitdefender or Norton to cover all your home's key devices with one subscription.

That leaves items such as your connected TV or smart devices like Alexa speaker or Ring doorbell. With these, you need to ensure they're set to automatically receive and install updates or take the time to perform a full device audit using Advanced IP Scanner (www.advanced-ip-scanner.com) or a similar tool, then check each device manually to see if updates exist.

You may also consider isolating some devices from the rest of your network. Check if your router supports VLAN—the Archer VR900 does, both wired and wireless. Synology routers should get this in the next update to the SRM operating system (1.3). If they are wireless devices, place them on the guest Wi-Fi network, making sure guests have no access to your main network. This may not be practical if these devices require access to local network resources to work properly.

If you'd like to split your network into 'trusted' and 'untrusted' zones, or want to create two separate, isolated networks, check to see if your router supports VLAN. Both Archer and Synology routers do, allowing you to separate wired from wireless, or create separate hybrid wired/Wi-Fi networks. Consult your manual for details on setting it up.

The final box reveals how to dial into your home securely using a

VPN server set up on your router, but what about more traditional VPN services, which allow you to mask your true location and encrypt data over untrusted Wi-Fi networks? Some routers, including the Synology, allow you to set up a VPN connection on the router itself, which means all the devices in your home can benefit without being connected separately.

On the Synology router, first check that your VPN provider supports L2TP/IPsec or OpenVPN (via a configuration file) for the most secure connection. Then navigate to Network Center → Internet and click 'VPN Settings' underneath 'Primary Interface' to set up the connection. Remember, though, that VPNs carry a performance penalty, so you might prefer to stick to protecting individual devices instead of slowing down your entire internet connection.



VLAN support is built into the Archer VR900 series

gamesradar+

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



New Channels: Movies, TV and Cool Stuff!


www.gamesradar.com

mobilism.org

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH



ACTIVATING THE GOD MENU

If you like to keep things all in one place, the 'secret' God Mode menu is a great place to organize important settings and tweaks in one folder. Simply head over to your desktop, right-click and create a new folder. Rename this GodMode. {ED7BA470-8E54-465E-825C-99712043E01C} and the quick-access control panel folder will give you over 300 options to play around with, which should save you some time.

MAKE - USE - CREATE



54
Design textures and patterns in Adobe Illustrator



60
Build a better file server with the Raspberry Pi



62
Optimize your Steam and Proton gaming on Linux



SAM LEWIS
STAFF WRITER

A LUCKY PURCHASE

In the last issue, I dived into the world of car coding by picking up a Carly OBD reader. After checking my own car, I was reassured to find that everything was as it should be. Luckily, I had opted for the annual subscription package that makes my code reader work with other car brands, as just a week later, my mother's car had a light flag up on the dash. To save her from going to a mechanic, I stepped in and diagnosed it within five minutes. Not only was it a quick alternative, but it also gave me the exact engine error code. I searched this online and within a week we had booked it in and the problem was sorted. Before anyone asks, no I didn't charge her for this service, I'm a nice son.

At the time of writing, we are a few days away from the release of *GTA: The Trilogy Definitive Edition*. I am super-excited, as I was too young to play the three 3D *GTA 3* releases when they originally came out. In the years since, I have played these titles on Android, but now is the time to play them properly. I'm glad they have kept the same cartoonish classic *GTA* styling as the old titles. The gameplay will hopefully be a rewarding and smooth experience, especially with *GTA 5* controls being included. I'll tune in next issue, hopefully with some good things to say about the game.

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



presents:

AUTOPSY

THIS MONTH WE DISSECT...

Apple MacBook Pro 16-inch



Each MacBook Pro includes three USB-C ports, an HDMI port, a headphone jack, and an even rarer SD card reader—not to mention the magnificent MagSafe 3 port.



This MacBook may look angry, but really it's just disappointed that the pull-tabs took so long to remove.

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.



Forget about RAM and SSD upgrades. It's all soldered on



BACKGROUND

It's Apple again, but the new MacBook Pro with its M1-derived uberchip. It's a monster performer, but will it come apart?

MAJOR TECH SPECS

- M1 Pro (APL1103) SoC with 10-core CPU, 16-core GPU
- 16GB LPDDR5 RAM (200GB/s bandwidth) and 256GB SSD
- 16.2-inch (3456 x 2234 pixels) Liquid Retina XDR LCD display, 1,600 nits peak, mini-LED backlight
- 100 Wh battery, 140W USB-C PSU with MagSafe 3
- Hardware-accelerated H.264, HEVC, ProRes, and ProRes RAW
- 3x Thunderbolt 4, HDMI, SDXC card, high-impedance headphone jack
- Force Touch trackpad, 1080p FaceTime HD camera

KEY FINDINGS

- Just look at those ports. With so many places to plug things in and so many dongles scorned, Jony Ive must be rolling in his ... Ferrari. MagSafe also makes a triumphant return to the MacBook, despite its iPhone moonlighting—good news for all the clumsy cord-kickers, the pet owners, and almost everyone. But we didn't come here to ogle the design. They say true beauty is on the inside, so let's see if it's true.
- Can you hear angelic music, or is it just us? Sure, the rear case is still secured by pentalobe screws, but there's no adhesive this far in—and the clips are even less terrifyingly tight. Compared with the 2019 15-inch Touch Bar MacBook Pro, these models make even better use of interior space. Speakers and batteries fill every cranny, without sacrificing a robust-looking cooling assembly. Differences between the 14-inch and 16-inch models are rare: rear cover clips, chip arrangements, and hinge bracket shapes.
- The four outer battery cells all sport subtle but noticeable pull tabs, aka stretch-release adhesive. Pull these things to stretch out the adhesive, and, in theory, whatever's attached to it falls right out. Even better, it appears this battery isn't trapped under the logic board. That could enable battery swaps without removing all the brains first. But we don't see any pull tabs on the two center cells, and they refuse to budge. Are we screwed—or, worse, glued? This stumps us, until we had the brilliant idea to remove the trackpad. There are pull strips under the remaining battery cells, accessed through cutouts in the chassis, beneath the trackpad.
- The logic board—and accompanying heatsink and fans—are beefy, and this isn't even the M1 Max model. The M1 Pro is flanked by two memory modules (four on the M1 Max), soldered to the SoC for maximum bandwidth. These memory modules make up Apple's 'unified' memory, meaning they can be shared between the CPU and GPU, dramatically increasing the pool that both can draw from. What does all this fancy-pants unification mean for reparability? Everything is on the board. Future-upgradeable RAM? Nope. Better shell out as much as you can afford up-front. Future-upgradeable storage: technically possible, but highly impractical.
- Repairability Score: 4 out of 10 (10 is easiest to repair). Apple's M1 tech is rocking the industry, but its reparability isn't advancing as quickly. Still, this design represents a move in the right direction. The battery replacement procedure is improved. May we never again relive the nightmare of surgically removing cemented-in batteries from a "Pro" laptop. We've also griped about soldered memory in the past. We still don't love it, but it's hard to argue with the performance and battery life advances. The biggest blight? Soldered-down, non-removable storage—a major barrier for reparability, upgradeability, security, data recovery, and overall flexibility. 🛠

Textures, & repeated patterns in Illustrator

YOU'LL NEED THIS

ADOBE ILLUSTRATOR
CC 2020

<https://www.adobe.com>

WE HAVE ALREADY CREATED some interesting shapes, designs, and flat artwork previously in these tutorials, but there are always more steps you can take. As long as your creativity is infinite, you can add more to your design work and take it to the next level. Patterns and textures are a great route to advance your artwork. Sure, the less-is-more technique is also another strong approach, however, we are going all out for this tutorial.

Adding a texture can make certain elements in your work pop, giving them stronger contrast and a 3D look. When used in the right places, it gives more depth and detail that should help to grab the audience's attention. With flat artwork, a texture can give a grainy vintage feel, which is the look we are aiming for here.

In this tutorial, we will take a look at how to change the original flat art plant designs we created in an earlier tutorial and spice them up a little. After that, we will look at how to repeat the design to get a pattern and create a wallpaper out of this art. —SAM LEWIS



1 STARTING OFF

To start things off, create a flat piece of artwork onto which you can then apply this textured style. It's always best to make a duplicate file so you never overwrite the original—after all, keeping organized is key to good creative practice. Once you have duplicated the original file, name it appropriately to suit the style and now you can get going.

» As always, you should remember to set the color mode to suit the type of media that you intend your design to be viewed in. If it's physical media, make sure it's CMYK and if it's digital media, then check you're working in RGB.

» To create these textured effects, we will be using grain brushes, specifically stipple brushes. You can find packs of these brushes online for free or as part of a paid membership for Illustrator packs. There are tons to choose from but when you find one that you are happy with, download its AI file, open it up then copy and paste the brushes into your main art file.

» Once done, these brushes will appear in your brushes windows of the Illustrator file and you can use them on strokes and when using the brush tool. [Image A]

2 JUST A LITTLE STIPPLE

Now that we have our artwork sorted and the brushes ready, we can begin by adding the textures. Stipple brushes are great for adding the vintage or gritty look that can really transform flat artwork, so will be perfect for this tutorial.

» To make more use of the textured effect, we added some more elements to our original artwork, including some more plants, a shelf, and a mirror.

» The first thing to do is to click on the shape to which you wish to add the effects and add a stroke to it. As shown in our example, we already have some shading on the plant pots. We took this shading, added a white stroke to it, and then applied one of the stipple texture brushes to the stroke to get the desired textured look [Image B].

» To alter the strength of the texture, we decreased the size of the stroke to 0.75pt, as seen in our image. Even this slight change gives the shadows a deeper and grittier look.

3 ADD SOME TEXTURE

Next, we apply this to the leaves of the plants, giving them a cartoonish pattern that will hopefully make them stand out more. Instead of adding this texture to the outside stroke, we will be adding it into the shape [Image C]. Whenever we do this to the inside of the shape, we need to duplicate the base shape first.

» To make things easier, create a new layer where the duplicates will go. Then hold shift and select the shapes you want to use. Copy these shapes using CTRL + C, head over into the new layer, and paste them in the same place using CTRL + F.





C

» After this, use the brush tool to draw an outline slightly around the shape you want to texture with one of the stipple brushes. Take this path and send it to the back. You can do this by right-clicking on the brush line, hovering over arrange then clicking send to back. [Image D]

» Now use the direct selection tool to select both the original shape you traced around and the brush layer that you just sent to the background. With both of these selected, press CTRL + 7 to clip these. This will cut out all the brush that exceeds the original shape, therefore only putting the texture inside.

» As we added texture to the dark part of the leaf, we chose the lighter color of the plant to use as the brush color. This created a nice contrast and kept to the original color scheme.

4 MIRROR, MIRROR

Hopefully, the plants and their pots should now be finished with a nice bit of texture added to them, but what about the rest of the image?

» First, holding the shift key, we draw a straight horizontal line under the shelf and use a stipple brush effect to create a textured shadow. We decrease this to 0.5pt and use the same color as the shelf. For the mirror, we create a rough outline of the circle and apply a brush to that stroke too. However, this time we increase the point size to 5pt, which gives us a dispersed splatter effect to sit around the mirror. We sent that behind the mirror as we didn't want the pattern inside.

» Finally, using the pen tool, we add a curved line towards the top left of the mirror to replicate a glass effect. We then changed the stroke to a brush stipple effect to fit the rest of the artwork.

» So there you have it, a quick and easy method of creating more texture with the handy help of the brush tool [Image E]. This technique can be applied to a variety of media and can be used as the finishing touches to artwork as we have shown here.

REPEATED PATTERN



Now that you have completed this first part of this tutorial, you can experiment with ways of taking the concept further. One possibility could be to take one of the plants from the artwork we have made and apply a grid repeat effect so that you end up with a patterned wallpaper look.

First, we need to copy one whole plant by highlighting the necessary shapes—as an aside here, if you ever highlight too much of the board or too many individual shapes, simply hold the shift button and click on the parts that you don't want to highlight, and this will deselect them.

Once you have the desired items selected, make sure all the shapes that make up that item are grouped together, by pressing CTRL + G. Then, head over to the top object menu. Inside this drop-down, find the repeat menu then inside that, click the grid option.

As you can see, this has created a partial grid pattern of the design. Adjust this until you get your desired pattern. There are plenty of adjustments that you can make to change this grid. Firstly, you can resize the box—we extended it to fit our artwork. Then using the two sliders along the top and left-hand side, you can alter the vertical and horizontal gap sizes between the pattern. Another tweak we made is to resize the original object, this affects the whole pattern but, thankfully, you can see an outlined preview of what this looks like when changing it.

After playing around with the grid, we ended up with an even pattern that filled our artwork. You can use this on a range of different designs to create a pattern. This is a more accurate method that beats simply copying and pasting the original object over and over again.



D



E

Upgrade Raspberry Pi to Bullseye

YOU'LL NEED THIS

A RASPBERRY PI 3 OR 4

An internet connection
Spare SD card

A NEW VERSION of the Raspberry Pi OS was released not too long ago, and while this may elicit the same groans as Windows 11 did, there's actually a good reason to upgrade—if you have the right chipset. Not all Raspberry Pi 4 boards, it seems, are created equal. We already knew they differ in terms of the amount of RAM that's attached, but it seems that if your Pi 4 has the newer BCM2711 SoC with a C0T model number—the chipset introduced with the Pi 400—then you get a 300MHz speed boost from 1.5GHz to 1.8. You may well have already overclocked your Pi, but it's welcome nonetheless.

The new OS is based on Debian 11 Bullseye [Debian releases are named after characters from the *Toy Story* film and Bullseye is Woody's horse] and comes in 32-bit and 64-bit varieties. We're concentrating on the 32-bit release here, but if you want to access all the RAM on your Pi 4 8GB, you're going to want the 64-bit release.

You can either upgrade your Pi's everyday OS or put the new one on a clean SD card (or USB SSD) to try it out before switching over. You should note that this OS is new and may be a little flaky in places, but here's how it's done. —IAN EVENEDEN

1 SO WHAT'S NEW?

It's not exactly a Windows 11-style visual reinvention for Raspberry Pi OS, but there's a new window manager, Mutter, in town—unless you're using a 2GB (or less) Pi, in which case you're stuck with the old openbox, as Mutter writes to memory before drawing to the screen, and 2GB isn't enough for it. The UI toolkit has changed too, from GTK+2 to GTK+3, which might not sound like much of a change, but is the reason this update has been so long in arriving. The older version of the OS relied on GTK+2 features that aren't replicated in GTK+3.


2 UPDATE FIRMWARE AND OS

It's worth checking if your firmware is up to date, so type "sudo rpi-update" into a Terminal window and type "y" when prompted. If an update is applied, you'll need to reboot. Then run "sudo apt update" and "sudo apt dist-upgrade" to make sure your existing OS is up-to-date and ready for the upgrade. It's also worth noting that Raspberry Pi itself does not recommend an in-place upgrade like this. It's bound to break stuff and neither they nor we take any responsibility for when it inevitably all goes wrong. You have been warned! Back up your SD card now—use the Pi's SD card copier app to clone your card.

NEW THINGS

Apart from a new window manager and visual improvements, you benefit from running a newer Linux kernel with better security and more efficiency. If you're running the right Pi, you get a 300MHz speed bump. There's a move toward using more open-source software, with the camera module being operated by `libcamera` rather than `raspistill`. Compatibility with every module and HAT isn't there yet, but the new version of Chromium brings hardware video acceleration and makes watching YouTube on the Pi less of a slideshow.

There's now a notification manager, so if you pull out a USB drive without ejecting it properly you'll be told off in the form of a gray rectangle at the top of the screen.



3 EDIT SOURCES


The Sources file tells the OS where to look for updates. At the moment it's searching a folder called Buster on the Raspberry Pi repository, as that's the name of the older OS. To get the new one, open Sources.list in the nano text editor with "sudo nano /etc/apt/sources.list" and look for a line that says "deb <http://raspbian.raspberrypi.org/raspbian/> buster main contrib non-free rpi". Edit 'buster' to 'bullseye' and leave the rest the same [Image A]. Type "CTRL + x" to exit, "y" when prompted to save, and "Enter" to overwrite the file.

4 ANOTHER UPDATE

Now we're looking in the Bullseye directory, we can run the update with "sudo apt update" in Terminal [Image B] followed by "sudo apt dist-upgrade" and "y" when prompted. You'll see a lot of text scroll up the screen, and it will take some time.

5 AN ERROR?

Oh yeah. We got errors. A message reading "Some packages could not be installed. This may mean that you have requested an impossible situation... The following packages have unmet dependencies..." popped up and stalled the upgrade process. We asked an expert for some advice and eventually cleared it by



installing a new GCC compiler with “`sudo apt install gcc-8-base`”. Once it’s finished installing, type “`sudo apt dist-upgrade`” again and the process should complete. Or maybe it won’t, that’s the joy of early-version open-source software.

6 IT WORKED, KINDA

A final restart brings up a new-looking desktop. It has the same wallpaper and folders, but the trashcan and Documents folder icons are different, the Raspberry menu icon is a green arrow and the mouse pointer is invisible. We know it’s there, as right-clicking brings up the context-sensitive menu on the screen. We just can’t see it. Connecting via VNC (the settings weren’t wiped by the upgrade and we jack straight in from our Windows PC) solves that problem. The pointer is now a little ‘x’, but at least we can see it. A larger problem is that window headers aren’t drawing, meaning we can’t drag them around or minimize them, and have to use the File menu to close them. Once again, new software. Then the menus stop vanishing after we’ve moved on to the next, leaving a big mess on the screen [Image C], and we give up and try a fresh installation.

7 A FRESH INSTALLATION

Download the OS from <https://www.raspberrypi.com/software/operating-systems/> and the Imager from <https://www.raspberrypi.com/software/>. We chose the 32-bit full OS, a 3GB zip that expands into a 9GB ISO, but smaller installations are available. We installed it onto a 64GB micro SD card [Image D], but you could get away with a 32GB one.

8 IT’S ALIVE!

What we got was something that looked and acted exactly like the old Buster release of the Raspberry Pi OS. The new-look trashcan and folders we’d seen when doing the in-place upgrade didn’t appear, and we had to resort to “`cat /etc/os-release`” in the Terminal to check that we were, indeed, running the Bullseye release.



OVERCLOCKING A RASPBERRY PI 4

While the 300MHz speed boost from the new OS is welcome, the CPU in the Pi 4 can be pushed even harder if you stick a fan on it for active cooling. Kits are available online, some even have RGB! Overclocking is a simple matter of editing a text file. You’ll need the latest firmware and a bit of luck. We’ve taken a Raspberry Pi 4 with a heatsink and fan to 2,200MHz, but have heard of others that won’t go past 2,000.

Once fully updated, open Terminal and type “`sudo nano /boot/config.txt`”. Find the commented line about `arm_freq`, uncomment it, and edit it to read “`arm_freq=2000`”. Then add three more lines: “`gpu_freq=750`”; “`over_voltage=6`”; “`force_turbo=1`”. Note that setting `force_turbo` may void your warranty. Save and restart. Sometimes a full power cycle is necessary rather than a software restart.

Use “`watch -n 1 vcgencmd measure_clock arm`”, and your CPU should be buzzing away at 2,000MHz all the time. You can nudge the figure upwards, but it generally tops out at 2,200MHz. It’s possible to overclock other Pi models too, but with less speed and voltage. Tables are available online.

9 CHECK YOUR CLOCK SPEED

The line “`watch -n 1 vcgencmd measure_clock arm`” will give you a readout of the current CPU speed on your Pi. If you’ve got one of the models with the BCM2711 chipset, you can use this to check if you have the free speed boost by typing it into Terminal then opening a web browser. By default, the Pi CPU idles at 700MHz but speeds up when needed while remaining passively cooled.

10 DOES IT WORK?

Yes, it does—at least some of the time. However, it would probably be better to class this release as almost beta software at the moment, and we’d recommend sticking with the Buster release wherever it’s project-critical. The ability to keep different OSes on different SD cards makes this easy. As it stands, there are some compatibility issues that will need to be worked out, and the new Chromium (version 92) crashed every time we tried to open YouTube. The first page it loads on opening is a Raspberry Pi page encouraging you to buy a Raspberry Pi, which seems unnecessary, but other sites we visit every day open without complaint, and we are able to stream video from other sites. Who needs YouTube anyway? ☹️



Machine of the Month: Apple IIGS (1986)

YOU'LL NEED THIS

A COPY OF GSPLUS

IIGS system ROM
Bootable hard drive image
<https://apple2.gs/plus/>

THIS MONTH WE LOOK AT what was arguably one of Apple's best and most underrated products: the Apple IIGS. Designed for superior graphics and sound (hence the "GS" moniker), the IIGS brought the old Apple II line into the next era of computing and was the first Apple computer to feature a color desktop.

Despite its impressive 16-bit capabilities, the IIGS would play second fiddle to the likes of the Commodore Amiga or Atari ST. This makes it a gold mine for retro computing enthusiasts, with cult games, applications, and a phenomenal sound chip waiting to be discovered. —JOHN KNIGHT

1 DEVELOPMENT

Steve Wozniak rejoined Apple and set his sights on making a substantial upgrade to the aging Apple II line, with a 16-bit CPU and backward compatibility with 8-bit Apple II software.

» For a CPU, Wozniak chose the Western Digital 65C816: a supercharged 16-bit descendant of the Apple II's old 6502 processor, which would later be used in the Super Nintendo. This was combined with Apple's custom "Mega II" chip, which provided Apple IIe hardware emulation in a single chip, allowing '95 percent backward compatibility'.

» For graphics, Apple used a custom IIGS Video Graphics Chip (VGC), which was capable of 640x200 resolution, and 256 on-screen colors from a palette of 4,096 colors.

» For audio, Apple chose the highly advanced Ensoniq 5503 Digital Oscillator Chip, designed by Robert Yannes—the man behind the legendary SID chip in the Commodore 64. This monster chip was used in professional synthesizers, and when pushed to its maximum, allowed for 32 voices and eight stereo channels.

2 LAUNCH

Debuting in September 1986, the IIGS had a launch price of \$999 (or \$1498 with a monitor), which was much cheaper than the \$2,599 Macintosh, the computer upon which Apple was betting its future.

» In most areas, the IIGS was more powerful than the Macintosh, except for its CPU. Although the 65C816 was capable of speeds up to 14MHz, Apple deliberately underclocked the CPU at 2.8MHz so as not to compete with the Macintosh.

» For an operating system, the IIGS only launched with ProDOS 16—a more advanced 16-bit version of the DOS used by the Apple II—and a rudimentary (though stylish) graphical interface for browsing and launching system files.

» ProDOS 16 soon shipped with Apple II Desktop, a rebranded piece of third-party Apple II software that provided a simplified version of the Macintosh desktop. This was soon followed by Apple IIGS Finder, a stripped-down Macintosh-style interface in color, with proper file management.

» However, in 1988 the IIGS received a huge improvement with GS/OS. This was a proper desktop, closely resembling the interface from the Lisa and original Macintosh, except GS/OS was in full color and had greater cross-platform compatibility.

3 SOFTWARE

Like the Macintosh, many IIGS games are simple affairs. However, there are plenty of action and arcade games.

» Some standouts include *Arkanoid II* (1988) and *Rastan* (1990). *Quadronome* (1989) is a personal favorite, providing a cool

3D variation on *Pong*, and *Silpheed* (1988) also moves vertical shooters onto a 3D plane.

» Adventure fans should try *Neuromancer* (1989), with 1992's port of *Another World*, named *Out of This World*, possibly being the system's most impressive title.

» The IIGS excelled at desktop publishing and multimedia. *Paintworks Plus* (1986) provided a colorized equivalent of *MacPaint*. The later *Paintworks Gold* (1987), *Deluxe Paint II* (1987), and *Graphics Studio* (1988) provide better experiences. Animation fans should try out *Fantavision* (1987) and *Animasia 3-D* (1994).

» MIDI (musical instrument digital interface) is a real strength of the IIGS, and *synthLAB* (1989) is an easy way to experience IIGS sound. For something more advanced, check out *Music Studio 2.0* (1987), *Noisetraacker GS* (1992), and *SoundSmith* (1990). *Diversi-Tune* (1988) has cool bouncing ball karaoke.

4 LEGACY

The IIGS received updates each year, with more advanced releases of GS/OS, firmware updates, and increased RAM. By 1989, Apple was shipping with 1MB as standard. Third-party accelerator boards allowed owners to unlock the IIGS's real potential.

» The IIGS proved a moderate success but never received full backing from Apple, with staff actually leading customers and developers away from the IIGS and toward the Macintosh instead. Production ceased in December 1992.



» Fans of Commodore's SID chip will enjoy Bob Yannes's later work with the IIGS's powerful sound processing.



» If only the IIGS had received the company backing it deserved, Apple's fortunes may have been very different in the 1990s.

» We couldn't find any official sales figures, but somewhere between 1.25 and 2.6 million units were produced – if anyone has an accurate sales figure, let us know.

5 EMULATION

We found emulating the IIGS was easier than the Macintosh, but there are still some hoops to jump through.

» To get started, you will need an IIGS system ROM and life will be easier if you can also find yourself a hard drive image.

» ROM files are easy to find but illegal to download. However, you can dump your own ROM from a real IIGS. Instructions are available at macgui.com [which for legal reasons is what we'll assume you're doing].

» Hard drive images *should* be okay, given that Apple has a free download archive of "obsolete" software, but we can't guarantee anything. We recommend System Software 6.0.1 if you can get it—see the box for more information.

» Several IIGS emulators are available, but we had the best experience with GSplus, available for Windows, Linux, and Mac. Download links are provided, with Windows users given a Zip file and Linux users provided a .deb package and a binary tarball.

» For Linux users who have installed the .deb package, enter the command:

```
$ gsplus
```

» For Windows users, extract the Zip file, open the new folder in your file manager and launch the gsplus.exe executable.

6 FIRST RUN

The first time you start GSplus, you will encounter an error about a missing ROM file. Click OK to continue.

» GSplus will just display garbage until you supply a ROM image. Press F4 and you will be prompted with a ROM file

FINDING SOFTWARE

The IIGS isn't as well archived as the Apple II or Macintosh, but many of the same websites host IIGS software, such as macintoshrepository.org. As usual, archive.org has a good selection, but our favorite site was "What is the Apple IIGS?" at whatisthe2gs.apple2.org.za.

For hard drive images, archive.org has a hard drive image available in T0SEC's IIGS archive, though we preferred the "What is the Apple IIGS?" website, which has hard drive images pre-loaded with shareware and freeware software.

selector. Press Enter and a file browser will appear. Browse your file system by using the arrow keys and Enter. Once you have located your ROM file, press Enter.

» Once selected, you will be back at the ROM File Selection screen. Press down to select 'Back to Main Config' and press Enter.

» Now that you're at the main menu, open the Disk Configuration menu. The Disk Configuration screen looks intimidating at first, but don't panic. Floppy images are assigned to anything under s5 and s6, and hard drive images are assigned to anything under s7.

» Press Enter on s7d1 to choose your primary hard drive image. A file browser will open, where you can select your image in the same way as with the ROM file.

» If you have other hard drive images, these can be opened at the same time as your primary image by assigning them to s7d2, s7d3, and so on.

» Any successfully mounted disk images will appear on your desktop later. This applies equally to any floppy images, mounted under the s5 and s6 entries.

» Once finished, hold the down button until you reach 'Back to Main Config' and press Enter. Now you can either scroll down to the Exit Config menu entry or simply press F4. GSplus will automatically save your new configuration by default.

» Assuming you chose a working hard drive image, GSplus should now boot your desktop. If it doesn't, try another image, and make sure it has the operating system pre-installed.

» As for shortcut keys, F4 opens the configuration menu, and if you want to enter full screen, press F11. Right-clicking will alter the CPU speed. 🖱️

SPECIFICATIONS

APPLE IIGS (1986)

CPU:	Western Design Center W65C816 @ 2.8MHz
RAM:	256KB - 1MB (expandable to 8MB)
ROM:	128KB
Graphics:	Apple VGC, maximum 620x200 resolution, 256 colors on-screen from 4,096 color palette
Sound:	Ensoniq 5503 Digital Oscillator Chip, 64KB dedicated RAM
Storage:	5.25" and 3.5" floppy, hard disk and CD-ROM expansion
OS:	ProDOS 16, GS/OS
Released:	September 1986
Production:	1986-1992



» The more powerful IIGS machine allowed colored equivalents of classic Macintosh software, such as *Shufflepuck Café*, *Dark Castle*, and *Paintworks Plus*.

Building a better file server with the Pi

YOU'LL NEED THIS

RASPBERRY PI 3 OR 4

NextCloudPi
USB 3.0 HDD
Internet connection

RUNNING YOUR OWN CLOUD STORAGE SERVER saves money, allows you to expand storage as necessary, and can be done with a device as small as a Raspberry Pi. In our previous guides on the Raspberry Pi, we have covered everything you need to know about installing and setting up NextCloudPi and the Nextcloud Ubuntu Appliance, as well as the manual installation of Nextcloud on Raspberry Pi OS.

But even when the software is installed and connected to your network, the job is still only half-finished. You will need to configure it for your specific use-case, enable external storage to reduce wear and tear on the Pi's SD card and make the Nextcloud server accessible from beyond your home network.

The most straightforward way to set up external access to Nextcloud and to add external storage is to do it with the NextCloudPi installation. —CHRISTIAN CAWLEY



» Connect a USB drive or full-size external hard disk drive to your Raspberry Pi for storage with Nextcloud—USB 3.0 preferred!

EXTERNAL STORAGE

The Raspberry Pi 3 and Pi 4 boot from the SD card. While this can be changed in favor of a more reliable USB device, SD is the default option. This has obvious advantages, from speed and convenience to simplicity. However, there is a key disadvantage, which is the longevity of the card: regular reading and writing of data to the SD card will shorten its lifespan.

» Throw into the mix the fact that your NextCloudPi storage installation is limited to the size of the microSD card in your Raspberry Pi and you have an obvious problem. It's just not big enough to compete with Dropbox or OneDrive or Google Drive, and you're limited to the maximum supported capacity of 1TB.

» Adding an external disk drive to Nextcloud requires only a device with USB connectivity (unless you're using a case with a SATA adapter). For faster data transfer, use USB 3.0+ devices, and for the optimum, rely on SSD or NVME/M.2 flash.

» First, connect the USB device and then set up an SSH terminal to your Nextcloud and enter

```
sudo mkdir /media/USBstorage
```

» In your computer's web browser visit the NextCloudPi web panel (using the 4443 port, so [https://\[IP ADDRESS\]:4443](https://[IP ADDRESS]:4443)). If it doesn't launch automatically, start the wizard by clicking the wand icon in the toolbar. Select USB Configuration, input Yes, then Continue, and Format USB.

» When this completes, select 'Move data to USB'. You now have a Raspberry Pi-powered Nextcloud server that can store as much data as you need without wearing out the microSD card.

2 STATIC ACCESS

Configuring Nextcloud to be accessed from devices beyond your network (such as your mobile device while shopping) is possible with a static IP address. But before setting this up for external access, it's smart to configure the Raspberry Pi with a static IP address for access within your network and, more crucially, your router first.

» To confirm the Pi's IP address and how it was assigned, open a terminal or SSH connection and enter `ip r | grep default`

» If you see the letters `dhcp` the IP address is assigned dynamically by the router. Make a note of the IP address, then check the nameserver or DNS IP:

```
sudo nano /etc/resolv.conf
```

» Make a note of the IP address listed here for the nameserver before proceeding. Setting a static IP address can be achieved by editing the `dhcpcd.conf` file in `nano`:

```
sudo nano /etc/dhcpcd.conf
```

» Add the following lines to the end of the file, changing the capitalized text with the information you have to hand:

```
interface NETWORK
static ip_address=STATIC_IP/24
static routers=ROUTER_IP
static domain_name_servers=DNS_IP
```

» Change `NETWORK` to `eth0` or `wlan0`, depending on the network type (Ethernet or Wi-Fi).

» For `STATIC_IP` enter the IP address you wish to use for the Pi. It's smart to use the one already assigned to maintain connectivity.

» For `ROUTER_IP`, use the gateway IP for the router on your network. This is usually printed on the router.

» Finally, set `DNS_IP` to the IP address you made a note of after checking the `resolv.conf` file. Finish by rebooting with

```
sudo reboot
```



» **Set up cloud storage that you can access across the internet by configuring port forwarding via static IP or dynamic DNS.**

» Confirm the changes with Ctrl+X, then Y and Enter, and reboot. You can check the static IP address has been applied using `ip r | grep default`

» If you can't get a static IP address from your ISP, and many of them don't provide one (see boxout below), your best option for accessing the Nextcloud from anywhere is with a dynamic DNS provider.

3 ACCESS ANYWHERE

Start by accessing your NextCloudPi web panel and enter the credentials you used during setup.

» In the web panel, find Networking and select 'letsencrypt' from the menu. Add a working email address when prompted and click the wand icon to launch the setup wizard and select 'External access'.

» Answer Yes, then wait for setup to complete. When prompted for Port forwarding, choose 'I will do it manually'.

» You will need to forward ports 443 and 80 (for HTTPS and HTTP) to the Raspberry Pi, targeting the router-assigned IP address for the device or its hostname (`nextcloudpi`).

» How you do this will differ from router to router. Many routers offer an online help tool, but you can also refer to the device's documentation to find out how to use port forwarding.

» In the next screen, sign up to FreeDNS when prompted, select Subdomains and choose a free domain and subdomain, then input your router's public IP address (check by searching Google for 'what is my IP') and Save.

» Next, choose Dynamic DNS and copy the Direct URL for your subdomain into a text file. Remove everything before (and including) the question mark, so that you are left with a long alphanumeric string.

» On the FreeDNS screen in NextCloudPi, complete the Domain field with the subdomain you chose, then paste the alphanumeric hash from the text file into the Update Hash field and click Finish to complete.

» With your Nextcloud installation now fully online, it should be accessible from the mobile app or any other device through your browser.

» Note that by establishing a static IP, you may need to create a new connection in whichever version of the Nextcloud app you are using.

» Almost all these can be found in the apps list. There are several apps that can really enhance your Raspberry Pi cloud server, especially if you're planning on using it for a home file and media server.

» In the Featured Apps list, you'll find a list of snaps that are listed by type. In the Multimedia tab, you'll find apps for playing media.

» There is also an app called Automated Media Conversion, which converts media files using FFmpeg based on rules that you set. It can also convert files as they're uploaded.

» Antivirus and anti-ransomware apps can also be installed to protect your cloud data, along with integrations with mainstream cloud solutions.

» This is useful for easily migrating your cloud storage across rather than downloading to a PC and then uploading to the Nextcloud server.

CAN YOU GET A STATIC IP?

Internally, a static IP is relatively easy to set up: simply instruct the device what its IP address is (see main article) and go from there.

Beyond your network, it's a lot trickier. Dynamic DNS tools are useful and have largely replaced the need for a static IP from your ISP—but what if that's the route you'd prefer to take? Dynamic DNS solutions can fail, after all, making a static IP address far more reliable. Relying on your router's external IP address is a risk, as you don't know how often your ISP will refresh it.

In the past, a static IP address could be expensive. These days, however, having a static IP address is more affordable. Some domestic ISPs will let you have a static IP address for a monthly fee, others with a one-off setup fee. Some still require you to switch to a business account, which comes with some shortcomings you might not be comfortable with.

The best way to find out whether you can get a static IP address and learn about any associated costs is to chat with your ISP. You might be pleasantly surprised at how little it can cost!

4 ENHANCE WITH APPS

With the Nextcloud server accessible from anywhere, there are a few things you can do to take it to the next level.

Steam and Proton gaming on Linux

YOU'LL NEED THIS

STEAM

Linux
Proton
Internet connection

IF YOU'RE AT ALL INTERESTED in gaming on PCs, you have probably already come across Steam, Valve's platform for distributing, updating and running games. Steam makes it possible to purchase a game, install it over the internet and then run it from the Steam interface.

That brings us back to the age-old gaming conundrum of support if you happen to be a Linux user, as not every PC game is designed to run on anything other than Microsoft Windows. That said, there are plenty of Steam games that will run on Linux, and quite often, a Windows Steam game can be convinced to run on Linux, even though some 'fettling' by the user may be required. This approach is officially supported by Steam using a system called Proton.

Getting all of these things running, and then possibly optimizing the results, is what we'll be looking at in this tutorial. —MICHAEL REED

1 STEAM POWERED

Steam itself is installed through a custom program called the Steam Installer. This makes sense because Steam updates itself and the games you install with it without relying on Linux's own update systems. These days, the installer is in the official repositories of many Linux distributions such as Ubuntu and Fedora and their derivatives. As Steam is proprietary software you may have to enable a specific repository, such as 'Multiverse' for Ubuntu or 'Nonfree' on Fedora. Having done this, you can, for example, install Steam on Ubuntu as simply as typing `sudo apt install steam-installer`.

» If you're running a DEB-based distribution, but you can't find the Steam installer in the official repositories, you can obtain the installer directly from the Steam website (steampowered.com) and install it with the usual `dpkg -i [Name of archive]` as the superuser.

» There is also another non-official way of installing Steam that all Linux users might find interesting, and that is installation via the Flatpak system. This offers a few advantages of its own in terms of privacy and sandboxing, as Steam is a system unto itself once it's on your system. A standard Steam installation is probably safe to use in the vast majority of cases, but if you are concerned about safety or if the other installation methods don't work, give the Flatpak variant a try (search for 'Steam' on <https://flathub.org>).

2 GAMES IN STEAM

Let's get started with an example that we can use to explore some of the features of Steam. Use the search feature at the top of the Store page to search for OpenTTD, a free management game, and click it to go to the product page for that game. Underneath the preview images and movies, we find a bar with some pertinent information. The first point of interest is the price, which in this case is 'Free' (yay!). There are also three icons that indicate that the game is compatible with Microsoft Windows, macOS, and Linux.


» That final icon, a piston arm over a wheel, indicates that the game has native support for typical desktop Linux distributions in addition to SteamOS, Valve's custom, game-orientated Linux distribution. Valve sells specialized devices that run SteamOS, including the recently announced (and rather tasty-looking) Steam Deck hybrid console. Developments like these are good news for gaming on desktop Linux because it incentivizes Valve to keep up the support.

» It's worth scrolling further down the product page to determine what the recommended specifications are and make sure that your machine meets them. In the case of OpenTTD, they are fairly light.

3 INSTALLATION OPTIONS

As the game is free, we could begin the installation procedure, but there are a few points to make about installation options in advance. It's quite common to end up with a system that has more than one hard drive installed and, for this reason, you might prefer to install games to a drive other than the system drive.

» To begin to configure this, the option that you're looking for is called Steam Library Folders, and it's located in Steam → Library → Steam Library Folders. From here, you can add an alternative installation folder located anywhere in the Linux filesystem that you have permission to access. The great thing about doing this is that the default game installation location is still available as an option in the installation dialogue. So, you can make decisions about installation location based on available hard drive space and performance



» Enabling Steam Play within the Steam settings dialog so that we can run supported and unsupported Windows games on Linux.

©STEAM

THE STEAM INTERFACE



1. MAIN MENU

Easily overlooked, this is the menu bar. It's mostly concerned with options to configure Steam itself.

2. STEAM SECTIONS

These tabs take you to different parts of Steam. Each entry is also a pulldown menu if you hover over it.

3. GAME LIBRARY CONTROLS

Three filter icons to choose from; installed games, owned games and native Linux games.

4. SEARCH

The search box. If you need this, you probably have too many games!

5. GAME LIBRARY

The list of games in your library. Left-click to be taken to the game's page. Right-click for the properties menu.

6. RUN A GAME

A button that changes between play, install, and stream functions.

7. GUILT-O-METER

Total playing time. Keep this covered when your boss, partner, or parents are present.

requirements, particularly useful if you're rocking an SSD as the system drive. Note that you can install to other types of storage, such as network drives or USB flash drives, but the performance tends to be poor in the case of most games.

» To carry out the installation, click on 'Install' or (confusingly, in the case of free games) 'Play', near the top of the product page. Complete a test of the system by clicking on 'Play' once the installation has completed to check that the game runs properly.

» As we said, Steam handles its own updates, both for itself and the games that you have used it to install. The updates for games are often huge, but you can monitor, pause or cancel them on the Downloads page (hover over Library and click Downloads). This is fortunate because, as well as consuming bandwidth, the updates are compressed and encrypted and this means that they use a lot of CPU time while they are downloading. In our experience, they can keep two cores fully occupied for the entire duration. Handily, you can exit Steam at any time, and updates will resume next time it is launched.

4 WINDOWS GAMES

But what if you want to run a game that doesn't offer a Linux version? In such cases, there is a good chance that it can be made to run using a system called Proton. An official part of Steam on Linux, Proton makes use of both Wine and DXVK, layers that translate Windows software and graphics calls into native Linux ones.

» However, there are some caveats. Firstly, compared with how the game runs natively on Windows, performance and stability might be an issue.

» Secondly, it raises the ethical question that, arguably, running the Windows version under Linux might reduce pressure on the developers to make a Linux native version of a game.

» On the first point, the best way to get answers about the viability of running a particular game is to search for it on the ProtonDB website (www.protondb.com) to see



» Getting ready to play Firewatch on a laptop without 3D support, thanks to Steam Remote Play streaming.



» Examining the compatibility reports for *Metro 2033* to work on ProtonDB. It's a pretty detailed report and looks promising.

what luck other Linux users had with it. The database itself is well organized and includes information about when the testing was carried out and what the specific hardware and software configuration used was.

» The level of success of running the game on Linux is rated from 'Borked' to 'Platinum'. So, if you find a recent, highly encouraging report about the game running well on a similar set-up to your own, the odds are good. In addition, ProtonDB reports contain information on tweaks you can use to make stubborn games run properly. More on that later.

» To run Windows versions of games on Linux, visit the Steam Play settings (Steam → Settings → Steam Play). There are two options here to be enabled: "Enable Steam Play for supported titles" enables approved games to run. These are games that Valve has tested and knows to work properly.

» Ticking 'Enable Steam Play for other titles' enables you to attempt to run all Windows games that you have access to through Steam, even if Valve has not tested the game. Again, examining the entry for the game on ProtonDB will give you the best idea of what to realistically expect.

5 TROUBLESHOOTING AND TWEAKING

If you look through ProtonDB, you'll notice plenty of references to Winetricks (<https://wiki.winehq.org/Winetricks>). So what is it? Winetricks is a helper script that installs various Windows components that are needed by games.

The script pulls through the component and makes it visible to the Wine subsystem. It is worth knowing about, even if you're using Wine outside of Steam. For example, if you typed `winetricks corefonts` it would download the Microsoft Corefonts, a set of fonts often needed by Windows applications, and install them to a location that Wine, and therefore Windows applications, can see.

» Things get a bit more complicated if you want to apply Winetricks fixes to a specific Proton game, but not by much. Thankfully, there is a helper script for the helper script called Protontricks (<https://github.com/Matoking/protontricks>) for just this purpose. If you have a full Python set-up (see your distribution documentation to install Python) you should be able to install it with `pipx install protontricks`.

» To apply Protontricks to a specific game, you need to discover the specific game ID of that game, using Protontricks itself. For example, typing `protontricks -s fallout` found *Fallout* and *Fallout 4* installed on our Linux gaming machine. This might seem like a lot of work, but we have to admit that being able to add all of the resources that each game might need on a game-by-game basis is pretty cool.

» This is necessary as different games might need different versions of a resource. You might even find that this flexibility makes some older games more viable on Linux than on native Windows, and this will be increasingly so as time goes by. The syntax is `protontricks [game ID] [Winetricks command]`, but most of the time, you'll simply be cutting and pasting the needed commands from the ProtonDB entry.

» Moving on from Protontricks, some games need special launch options to be added. The dialogue for doing this is accessed by right-clicking the game name in the Library section in Steam, selecting 'General', and then entering the launch options that you have found on the relevant ProtonDB entry into the box. Most such options will end with the string `%command%` at the end.

» If you add more than one launch option at once, remove the extra instances of `%command%` as it should occur only once, at the end of the overall option string.

» If you have to alter the INI files of the game, you can browse to the game folder by right-clicking the game

RUNNING GAMES REMOTELY

Steam can run games remotely between two Steam-equipped setups. The way it works is that the remote computer runs the game and compresses the video and audio. It then streams this video and audio to the other computer while the client sends back mouse, keyboard, and controller input.

The two computers you use don't even have to be of the same architecture, and there's an Android app too. We tested it out between a Windows computer with a GNU/Linux one over a home network and found that it worked well.

The performance-limiting factors are the speed of the host that's running the game and the latency and bandwidth

of the network, rather than the power of the client computer. So you could have a tricked-out gaming computer with a decent graphics card hosting the game, paired with a low-powered Linux box as the client.

Over gigabit Ethernet, it was possible to run an action-orientated game with no noticeable input lag at all while running the game at 1080p. There was some slight degradation of graphics quality as video compression was in use. However, we were able to alleviate this to an extent by increasing the graphics balance setting (Steam → Settings → Remote Play → Client Streaming Options To This Computer: Beautiful).



» Streaming from a Windows PC onto a Linux PC. Hold Escape for the menu.

name, clicking 'Properties...' and then 'Browse...' in the 'Local Files' section.

» If any of that seems a bit confusing, don't worry; we're going to go through an example of getting an awkward game to run properly using these techniques.

6 AN APOCALYPTIC ADVENTURE

It's time to take a trip into the forbidden wastelands of lawlessness and chaos. No, we're not talking about rebooting into Windows. We're going to take a look at what it takes to get the post-apocalyptic, first-person, role-playing game *Fallout 4* running under Linux. Don't worry if you're not a fan of that particular game (heathen that you are) as we've chosen this one because it provides a good example of getting something a bit tricky to run properly.

» We had purchased the game and played it under Windows, and as it was already present in the Steam library, we installed it by selecting it and clicking the install button.

» Remember, as this game isn't approved by Valve, you have to enable installation in the Steam Play settings, as detailed earlier on. This led to a 39GB download that took ages and made the CPU fan spin quite fast while it was doing so. Steam and Proton added all of the components that the game needed to run, so we didn't have to make any use of Protontricks in this instance.

» At that point, all we had to do was to click on 'Play' to launch the game and it did launch as expected, but there were some problems. First, the sound was crackly and distorted for the introductory movies and entirely absent for the game itself.

» Second, the mouse kept moving in a crazy way, making the game unplayable as well as silent. Still, early signs looked promising for what is a technically demanding game. Leaving the game to find some fixes, we hit upon another problem, because it had hung rather than exiting cleanly.

» Taking the last thing first, the hang-on-exit problem was solved, in the short term, by process-killing the running game. We did this by pressing Alt+F3 to open a virtual terminal. From here, we typed `top` to view all of the user processes on the system in Top, seeing `Fallout4.exe` as the first item. `killall -9 Fallout4.exe` killed the process and pressing Alt+F7 got us back to the regular desktop.

» On subsequent runs of the game, we learned to close down cleanly by exiting to the main menu (rather than straight to the desktop) and quitting the game from there. Little workarounds like this are fairly common when you're running a Windows game under Proton.

» Searching ProtonDB, it turned out that the sound problems we encountered are common for *Fallout 4*. The solution was

GET THE BIG PICTURE



Steam can offer a gaming experience that is ergonomically closer to console gaming through what it calls Big Picture Mode. This pushes Steam into a full-screen mode with controller support. We have tried it out and we liked what it had to offer.

Big Picture Mode is invoked by clicking on a small icon at the top of the screen, and the first thing you'll notice is that every element of the user interface is jumbo-sized for better legibility on typical TV setups. You navigate around this interface using the mouse or a controller – all from the comfort of the sofa.

However, we found that the hardware support wasn't quite as slick as the interface undoubtedly is. For example, there is no clear way of specifying which screen should be used for the actual game playing. The most common workaround is to temporarily make the target screen your primary display, or set up display mirroring.

Similarly, there was some work needed to get the controller working properly in the games. You may have to experiment here to find a well-supported setup. Top marks for the interface, but room for improvement for ease of hardware configuration was our verdict.



» Download in progress. Keep an eye on transfers like these because they can consume a lot of bandwidth, CPU time and storage.

to cut and paste `WINEDLLOVERRIDES="xaudio2_7=n,b" PULSE_LATENCY_MSEC=90 %command%` into the launch options of the game. A quick jump back into the game confirmed that this had largely fixed the sound, but it took regressing to an earlier version of Proton, rather than the default experimental build, to make it work perfectly.

» The mouse problem was a little bit more complicated to track down. Sure enough, the suggestion, repeated over and over in ProtonDB, to add the line `bBackgroundMouse=1` to an INI file turned out to be the right one, but we initially had difficulty locating the file because it was duplicated in multiple locations, and it was impossible to figure out if what we were doing was having any effect.

» Eventually, we discovered that we needed to alter the `Fallout4.ini` file located in the `compdata` folder, within the simulated My Documents Proton folder, rather than the identical file within the game directory itself.

» At this point, we customized a character and ventured out into the apocalyptic wasteland, experiencing performance that was similar to running the same game on the same hardware via the Windows partition. Success! 🎉

LAB NOTES

JEREMY LAIRD, EDITOR



Running hot and heavy

Does power consumption matter?



Does Alder Lake have a drinking problem?

INTEL'S NEW Alder Lake chips are undoubtedly a big step forward. But do they have a fatal flaw? I speak, inevitably, of power consumption. Not in a global quasi-environmental sense, that's not my remit. What I'm talking about is the extent to which power consumption is a problem from a performance and usability perspective.

Last issue, we had the 12900K in the labs and, boy, does it like a drink. It's a 240W-plus chip. This month, the compact 12600K is in the house and, at first glance, it's a more efficient chip, topping out at around 150W in heavy multi-threaded workloads.

Depending on how you look at that, the 12600K seems to solve Alder Lake's power problem. After all, 150W is manageable on the desktop. On the other hand, the CPU part of Apple's new M1 Pro and Max chips delivers roughly the same class of multi-threaded performance for just 30W.

Yes, the 12600K is a fair bit faster than an M1, and it's a desktop chip, not a mobile chip. But the mobile version of Alder Lake is based on the same silicon. Clock it down for mobile use and it will use less power, but also deliver less performance. To put it another way, the 12600K uses four to five times more power for about 30 percent more performance.

Revised for mobile implementation, that gap will narrow. But it's hard to imagine that M1 won't remain at least twice as power-efficient, which will make Alder Lake look silly in a laptop performance comparison.

In the short term and for the desktop, you could say who cares? 150W is practical on the desktop and M1 chips only come with Apple computers. But it begs the question of whether the ARM instruction set, on which the M1 is based, is fundamentally superior to x86. People used to dismiss ARM as

not being suitable for high-performance computing. That now clearly isn't true. Instead, the current doubt is whether x86 can close the efficiency gap to ARM.

In the short run, it's academic but the comparison will become harder to ignore. Just imagine how much CPU performance 150W's worth of M1 could deliver on the desktop, let alone 240W. Anyway, M1 is showing what is possible when it comes to power and performance. We'll never have Apple chips in our PCs. We may never have ARM CPUs, either. But what we will have is definitely going to be faster thanks to the impressive example being set by Apple.



SAM LEWIS

Staff writer

Over the past year, I have fallen back in love with older games from my childhood and this nostalgia hit has left me with less time to focus on the newer releases. That was until November arrived.

Until then, I played the *Lego Star Wars* and *Batman* games on PS2 again. I'm also big into racing titles, so I always find myself hopping

back into *Gran Turismo 4*, *Dirt Rally*, and *Forza Horizon 4*. Other games I've gotten into include the original two *Star Wars Battlefront* games and a load of *Call of Duty* and *Battlefield* titles.

So when November crept up, I knew it would be a busy month. I was mainly interested in *Call of Duty: Vanguard*, *Forza Horizon 5*, *Grand Theft*

Auto: The Trilogy Definitive Edition, and *Battlefield 2042*. All of these games, on paper, I knew I'd be excited about.

The one I was most interested in, though, was *GTA Trilogy*. This might show my age, but when *San Andreas* came out I was only five, so the chances of my parents letting me play these titles were slim. I did manage to play them on

Android in my teens, but this should be an opportunity to play them properly before making a judgment. I will check in next month with my, hopefully positive, opinions.

Often with games, less can be more—a case of if it ain't broke, don't fix it, just tweak it a little. You don't need to reinvent the wheel, otherwise, it's back on with the PS2.

Editor's Pick: SkateBird

What do you get if you cross a chicken and a skateboard?



THIS IS one of the final pieces that the former editor Zak commissioned before he left, and it only exists because he found the concept so amusing.

There's a chicken, and it's on a skateboard, and it's called *SkateBird*. We're not dealing with high levels of sophisticated humor here, but that was never Zak's way.

However, I'm glad he did commission it, because I get to tell you that, for \$20 (it's on Game Pass too), you can have a lot of fun, and you can do it on integrated graphics. The recommended specs part of the game's Steam page doesn't mention graphics hardware beyond Intel HD 4000—the iGPU that came with Ivy Bridge chips almost ten years ago. Almost anything will be better than that, but this is a Switch game at heart and looks like one.

Grinding your way through bird-sized skateparks like Tony Hawk (that couldn't possibly be the inspiration, could it?) zipping up cardboard ramps, and grinding on drinking straws, the concept is clear. It's birds, on skateboards, and the more you skate the more birds are attracted to you. Being a bird, you can flap your wings for a sort of double-jump, but otherwise, you're basically doing ollies and hitting buttons on the pad to perform tricks. And I recommend you use a controller, as the kind of handling we're looking at here won't respond well to a keyboard.

The whole thing is remarkably easygoing. Fall and you get straight back up or warp to a reset point you chose with the touch of a button. The game isn't beyond the kind of cute spelling errors the internet is so fond of, either. Not only are you encouraged to 'Create-a-Birb' at the custom character creator (birds are bright



and colorful and wear a broad range of hats, it seems), but you'll also find yourself being asked to 'screem'—a squeak and flap of the wings—multiple times, as it counts as part of a trick and can be used to let out your anger when you come off.

The combo system fills a 'fancy' bar, which gifts you extra speed, and therefore height, once full. The birds just seem to be hanging on to the boards, never feeling like they're really in control. It's a wonderful kind of chaos, as eagles and larger species give you 'missions' that frequently involve collecting letters to make words.

However, it's also not without its problems. The camera regularly struggles to keep up with the action, there are timed levels where no timed level ought to be (in fact, timed levels should just not exist, period) and sometimes it can be incredibly strict on a point of trickery that you'd expect it to be as laissez-faire about as it is on others. The physics can be wobbly, in both entertaining and non-entertaining ways, and I'm pretty sure it's intentional.

The game may try to keep a straight face about things—there's a story about the lead bird's 'big friend' (its human owner) being sad because of his job, so the feathered pet gets to put on a bow tie, pirate hat, and glasses, to make him feel better—but it's almost impossible for the player to. There's bags of personality oozing out of *SkateBird*, and I'm glad that I don't have to put a score on the end of this, as it would need to be a 9 and a 4 at the same time, which is the kind of quantum superposition we could really do without.

So thanks, Zak, for finding avian skaters funny and thanks to Glass Bottom Games, a tiny indie team from Seattle, for coming up with the idea in the first place, even if it is a pun on tiny hawks. **-IE**
\$20, <https://skate.birb.rocks/>

Reviewed...



68 Intel Core i5-12600K



70 ASRock Radeon RX 6900 XT Formula



72 Dell S3222DGM

74 Eizo FlexScan EV3895



76 XPG Xenia 15 KC

78 AOC 24G2

80 MSI MEG Z690 Unify

83 Asus ROG Strix Go Core Moonlight White

85 SteelSeries Apex 3 TKL

87 Asus ROG Strix Impact II Moonlight White

89 AVerMedia PW513

90 Age of Empires IV

92 Google Chrome vs Brave





Intel's Core i5 CPUs tend to be great gaming chips...

Intel Core i5-12600K

Give it up for the new gaming champ

INTEL'S CORE I7 and i9 chips are all-round performance leaders, at least within Intel's own range of CPUs. But for gamers, it's actually the Core i5 that rules. Cheaper than an i7 or i9, i5 chips usually give up significant multi-threaded throughput but more or less match their more expensive siblings when it comes to in-game frame rates.

With the Core i5-12600K, the first Core i5 of the Alder Lake generation, that usual refrain may actually understate its abilities. That's because the 12600K is something even more impressive entirely. In many ways, it matches Intel's previous-gen top chip, even in multi-threading, but in a package that's almost half the price. That's Core i9 performance for Core i5 money.

The 12600K is a six-plus-four design. It has six Performance Cores (P-Cores) and four Efficient Cores (E-Cores). This is Intel's new hybrid approach to CPU design, which we reviewed last month in the top-end Core i9-12900K. The Performance Cores are codenamed Golden Cove and deliver maximum grunt at the highest possible clock speeds.

Meanwhile, the Efficient Cores are based on the Gracemont architecture, essentially a derivative of Intel's low-power Atom processor. They physically use much less die space and consume far less power. That does not, however, mean they are useless for serious computing tasks. Intel has said that Gracemont cores are roughly on a par with the old Skylake design of 2015 when it comes to single-threaded performance per clock.

Anyway, the total core count is higher than the old 11600K, which was a six-core chip, though it's worth noting that the 12600K's overall thread count comes to 16, since those Gracemont cores don't support HyperThreading. Put another way, with the 12600K, you're getting six high-performance cores, just like the 11600K but with higher per-core performance, plus those four high-efficiency cores. However you slice it, that looks like a big upgrade on paper.

As for the finer details, both P-Core and E-Core share access to 20MB of Intel Smart Cache (L3), with a maximum Turbo frequency of 4.9GHz for the P Cores and 3.6GHz for the E Cores. According to Intel's new power metrics, the 12600K's Base power is 125W with a Max Turbo power of 150W. The latter is much lower than the 241W maximum of the 12900K, immediately addressing that chip's most obvious flaw, excess power consumption.

All of the above is delivered via Intel's 10nm process tech, now rebranded Intel 7. Alder Lake is therefore the first 10nm Intel desktop chip family. But how good, exactly, is the 12600K at gaming? Absolutely outstanding is the answer. It blasts through our benchmarking suite, only just falling behind the much more expensive Core i9-12900K in a few titles.

More importantly, it squarely hammers the old Core i9-11900K flagship CPU in every game benchmark we ran. It's also far ahead of the AMD Ryzen 5 5600X, an important outcome for Intel's quest to regain all-round gaming supremacy. The one snag within this

story of untrammelled conquest involves *Assassin's Creed: Valhalla*. It simply will not boot on either of our Intel Alder Lake test systems, an ongoing issue that's likely related to the DRM issues Intel has officially confirmed with Alder Lake.

As for broader performance, the 12600K also crushes the Core i9-11900K in multi-threaded testing. The same goes for the Ryzen 5 5600X. Perhaps that's not surprising, given it has more physical cores than either. But it just shows those E Cores really do boost multi-threaded workloads. They're not just there to sip power, even if the 12600K beats the old 11900K while consuming far less juice.

Of course, some doubts remain over both the aforementioned DRM issues and the wider question of thread scheduling in Windows, plus the question of how performance is load-balanced across cores with hybrid architectures. We may yet find edge case scenarios where the hybrid approach is problematic. Those caveats aside, the 12600K jumps straight to the top of our gaming CPU leaderboard. Where the sheer thirst for power of the bigger 12900K makes us slightly uneasy, the 12600K is much easier to wholeheartedly recommend. It's a killer gaming CPU. —JACOB RIDLEY



Intel Core i9-12600K

■ A NEW HOPE Awesome gaming performance; much better power consumption than the 12900K.

■ PHANTOM MENACE DRM issues with some games; hybrid architecture somewhat unproven.

\$319, www.intel.com

BENCHMARKS

	Intel Core i9-12600K	AMD Ryzen 9 5900X
CineBench R23 Single (Index)	1,810	1505
CineBench R23 Multi (Index)	16,863	10404
X265 v5.0 (avg fps)	52	37
Total War: Three Kingdoms (DX11) @ 1080p (avg fps)	206	189
Far Cry 6 (DX12) @ 1080p (avg fps)	127	122
3D Mark: Timespy CPU Score	13,335	7786
Peak package power (watts)	113	76
Recommended Retail Price (\$)	\$289	\$299

Best scores in bold. Test bed consists of an Asus ROG Maximus Z690 Hero, 32GB (2x16GB) Corsair Dominator Platinum RGB DDR5 @ 5,200 MT/s, an Nvidia GeForce RTX 3080, 1TB WD Black SN850 PCIe 4.0 SSD, and an NZXT 850W PSU. All tests performed on Windows 11 on the highest preset available.

SPECIFICATIONS

Cores/Threads	6/12 P + 4/4 E
Base/Turbo Clock	3.7 / 4.9GHz
Architecture	Alder Lake
Lithography	Intel 7 (10nm SuperFin)
Memory Support	DDR5 @ 128GB (4800MT/s) / DDR4 @ 128GB (3200MT/s)
PCIe Support	PCIe 5.0 x16 (Graphics), PCIe 4.0 x4 (Storage)
Integrated Graphics	Intel UHD Graphics 770
TDP	125-150W

ASRock Radeon RX 6900 XT Formula

Big Navi revs its engine

THE RADEON RX 6900 XT currently reigns as AMD's most powerful GPU. ASRock takes that chip, slaps on a massive triple-slot, triple-fan cooler, and cranks the clocks to levels normally only seen on the smaller Navi 22 and Navi 23 chips. It's less about an elegant design and more of a brute force approach, but we appreciate the end result nonetheless.

If it's not immediately obvious, this card caters to extreme desktops with plenty of room. Perhaps there are some mini-ITX cases that could accommodate the card's girth, but we'd be concerned about limiting airflow. Also, finding a mini-ITX power supply with the required triple 8-pin connectors and probably 850W of output will prove almost impossible.

You could probably run the card off a 750W PSU, especially if it's a good quality model, but once you factor in CPU and motherboard power use, plus the 390-430W of power use (not counting for PSU inefficiency), and you can almost see the electrical grid flicker.

In the default mode, the Formula draws up to 390W of power, but who buys a card like this only to skip the extra factory overclock unlocked with ASRock's OC Tuner software? Kick the card into overdrive and it averaged just over 430W of power draw in testing.

Thankfully, the card runs relatively cool and quiet—and alternatively, you can drop power use and noise levels by enabling Silent mode in the software, but where's the fun in that? ASRock tuned the fans to keep the GPU at around 75C, and it has a noise profile of 48 dB(A) at 10cm. That's about the same as the reference 6900 XT. Enable OC mode and noise jumps to 52 dB(A), which is louder than AMD's reference design, but it also clocks in at nearly 2.6GHz—over 250MHz higher.

LIGHTS OUT

One strange thing about the card is that it actually looks quite boring, at least it does if you're into RGB bling. The only lighting on the card consists of a small strip around the back of the shroud, near the power connectors. ASRock's Phantom cards include RGB fans and more copious bling, and 'conservative' doesn't normally come to mind when we think about Formula One racing or muscle cars. But hey, it's what's under the hood that counts.

The performance ends up being about five percent faster than the reference RX 6900 XT, and it comes relatively close to matching the RTX 3090. It's actually faster than the 3090 in most of our test suite, but we also include *Control* and *Cyberpunk*

2077 running with ray-tracing enabled. That's the Achilles' Heel of AMD's RDNA 2 architecture, especially since both games include numerous ray-tracing effects. At 1440p, the ASRock card only manages to reach 43fps in *Control* and drops below 30fps in *Cyberpunk 2077*. The RTX 3090 is 60-80 percent faster in those two games, though *Assassin's Creed: Valhalla* favors the ASRock card (and AMD's GPUs in general) by 15-35 percent.

Our biggest complaint with the ASRock Formula is that it didn't go far enough. AMD released a liquid-cooled RX 6900 XT card to system integrators, and it includes 18Gbps GDDR6. We would have loved to see ASRock provide a similar VRAM for the Formula, as that could have boosted performance by another 10 percent in some cases. A decent overclock of the GPU can only go so far without a similar overclock to the memory subsystem, even if the Infinity Cache does help to mitigate that somewhat.

Our other major complaint is more a sign of the times than anything unusual. ASRock lists the MSRP for the RX 6900 XT Formula at \$1,899. It's high enough that you can find the card in stock for less than the MSRP, but that's because AMD's base MSRP is only \$999. Given the ASRock card achieves the same 65MH/s hash rate in Ethereum as other RX 6800, 6800 XT, and 6900 XT cards, miners aren't willing to pay the price premium. That means you could actually buy the Formula for around \$1,600. That's still more than we'd recommend, but these are desperate times for gamers. —JARRED WALTON

VERDICT



ASRock Radeon RX 6900 XT Formula

■ MUSCLE CARS Beefy cooling; extreme performance; 16GB VRAM.

■ ELECTRIC VEHICLES Expensive; inefficient; 16Gbps VRAM.

\$1,899, www.asrock.com

SPECIFICATIONS

Architecture	Navi 21
Lithography	TSMC N7
Boost Clock	2,475MHz
GPU Cores	5,120
Memory	16GB GDDR6
TFLOPS FP32	25.3
Bandwidth	512GBps
TDP	357W
Connectors	1x HDMI 2.1, 3x DisplayPort 1.4

BENCHMARKS

	RX 6900 XT ASRock	RTX 3090	RX 6800 XT
8 Game Average	105 / 65	112 / 74	94 / 58
Assassin's Creed Valhalla	120 / 70	90 / 60	108 / 63
Borderlands 3	147 / 81	128 / 76	133 / 73
Control (DXR)	68 / 43	109 / 71	62 / 39
Cyberpunk 2077 (DXR)	40 / 27	71 / 48	35 / 24
Dirt 5	148 / 97	149 / 100	139 / 90
Horizon Zero Dawn	145 / 82	137 / 94	136 / 76
Metro Exodus	113 / 72	115 / 77	99 / 63
Red Dead Redemption 2	123 / 83	117 / 83	108 / 72

Best scores are in bold. All testing conducted with a Core i9-9900K, MSI MEG Z390 ACE, 2x16GB DDR4-3600 CL16, 2TB XPG 8200 Pro M.2 SSD, Thermaltake Toughpower GF1 1000W. Scores are average framerates at 2560x1440/3840x2160, except in *Control* and *Cyberpunk*, which are at 1920x1080/2560x1440.



ASRock packs 4lbs of fun into its large and not particularly efficient 6900 XT Formula card

Dell S3222DGM

Affordable 32-inch, 1440p gaming

AT THE MORE extreme end of the monitor market, things don't just get a bit pricey. You can also find yourself with a display that's fine-tuned for some tasks, but not so good for others. Enter the new Dell S3222DGM. At just over \$300, it's an affordable 32-inch monitor and a decent proposition for everything from gaming to web browsing. It's a true all-rounder.

The S3222DGM is a 1440p model, which means 2,560 x 1,440 pixels. That's adequate in terms of desktop real estate, but also a good compromise between detail and performance. 4K is sweet, but it seriously hammers your GPU. Given the cost of graphics upgrades right now, that's a problem.

The issue with that 32-inch 1440p combo involves general productivity duties. The consequent 93dpi pixel density doesn't exactly make for sharp, clean fonts. Even photos and video streams look less crisp and detailed than we'd like, but let's park that while we consider the S3222DGM's wider qualities.

It's based around a VA-type panel with 165Hz refresh. The response time claims are 2ms gray-to-gray (GtG) and 1ms MPRT. Helpfully, Dell also rates each level of the user-configurable overdrive available in the OSD menu. 'Fast' mode equates to 8ms GtG, 'super fast' nets you 4ms GtG, 'extreme' is good for 2ms GtG, and 'MPRT' mode is rated at 1ms.

In practice, that latter mode hammers brightness so much we wouldn't touch it, while 'extreme' mode is a little overcooked with some visible overshoot. 'Super fast' is the best compromise with a good, but not stellar, response. Speaking of brightness and vibrancy, this is a purely SDR rather than HDR monitor, with a peak brightness rating of 350 nits. That's modest backlight power by today's standards, something that's obvious the moment you fire up the S3222DGM. There's nothing strictly wrong with the core image quality—no major flaws, no banding, no obvious compression—but the slight lack of zing and punch make this a fairly subdued monitor.

Still, the 165Hz refresh combines nicely with the decent response when it comes to gaming duties. For sure, if trigger-happy online gaming is your sole PC purpose, an even higher-refresh 1080p model will suit you better. If you're more into eye candy, drinking in the ray-traced visuals of *Cyberpunk 2077* or soaking

up the spectacle in the *Forza* series of racing games, the lack of overall fizz and low pixel density detract a little from the drama. But the S3222DGM still looks pretty good. For a multi-purpose gaming remit, the size, refresh rate, resolution, and response make for a sound all-round solution at this price point.

There are two HDMI 2.0 ports and a lone DisplayPort 1.2 socket, plus a slim-bezel design and a stand that offers height and tilt adjustment, but no swivel or rotate into portrait mode. At just over \$300, it isn't the deal of the century, but it is a convincing real-world proposition.

Did we mention that the S3222DGM is curved? It's such a gentle 1800R sweep that it's easy to forget it isn't flat. But at the same time, we're not convinced the curve adds much or that a 32-inch 16:9 aspect panel really needs any bend.

Our main reservation with the Dell S3222DGM remains that mediocre pixel density. It gives this monitor a clunky, coarse feel on the desktop. But then at this price, something's gotta give. And as an all-rounder with more emphasis on gaming than the boring stuff, Dell has got this one just about right. —JEREMY LAIRD

VERDICT Dell S3222DGM



TRANSATLANTIC SLEEPER SEAT

Good all-round feature set for the money; decent response and refresh.

LONG-HAUL CATTLE CLASS Subdued image quality; mediocre pixel density.

\$329, www.dell.com

SPECIFICATIONS

Panel size	32-inch
Panel technology	VA
Native resolution	2,560 x 1,440
Aspect Ratio	16:9
Refresh rate	165Hz
Response time	1ms MPRT, 2ms GtG
HDR	No
Contrast	3,000:1
Color	99 percent sRGB
Brightness	350 cd/m2
Video inputs	DisplayPort 1.2 x1, HDMI 2.0 x2
Other	AMD FreeSync Premium





Dell's 32-inch gaming panel combines 1440p and 165Hz refresh for a little over \$300



Eizo FlexScan EV3895

Epic but pricey productivity in a 38-inch package

THIS IS FAR from the first 38-inch nearly-4K ultrawide monitor on the market. But the Eizo FlexScan EV3895 is unusual in that it is aimed squarely at the professional productivity market, if not actually at content creation pros.

In fact, there are a few good reasons why we haven't seen a 38-inch, 21:9 curved productivity panel before. The curvature isn't necessarily a benefit when it comes to getting stuff done as opposed to, say, gaming immersion. Likewise, the 3,840 by 1,600 native resolution is a bit betwixt and between, falling short of full 4K and so limiting native workflows to 1080p, or at a pinch, 1440p.

The large panel size also translates into a relatively low pixel density, which isn't optimal for font rendering, general image sharpness, and detail.

A 'conventional' 16:9-aspect 32-inch or 27-inch 4K monitor offers more pixels, better pixel density, and the potential for more desktop space depending on your scaling preferences. It's also worth noting that the Eizo FlexScan EV3895 is a pure SDR rather than HDR, panel.

Of course, bona fide HDR content creation panels are specialized and hugely expensive devices, so the lack of entry-level HDR support probably won't be a deal-breaker for more generalized productivity. Anyway, the EV3895 conforms to Eizo's signature design and engineering ethos. In other words, it's sober-suited and light on visual thrills, but heavy on quality and practicality.

The fully adjustable stand feels robust and nicely engineered, and the chassis is a typically high-quality item. The height

adjustment is a particular highlight, delivering a range of 20cm via a clever and smoothly operating mechanism. The curvature of the panel is a gentle 2300R. Eizo rates the EV3895 at 1,000:1 for static contrast, 300 nits maximum panel brightness (there's no HDR support here), and 5ms response times. The panel itself is native 8-bit color with 14-bit LUT and what Eizo characterizes as maximum color space coverage of sRGB. DCI-P3 coverage is pegged at a claimed 94 percent.

Panel aside, the EV3895 offers excellent connectivity. The USB-C port combines DisplayPort Alt mode for full 60Hz refresh at native resolution with 85W of charging power and pass-through to both the integrated gigabit LAN port and three USB-A ports.



This 38-inch IPS beauty lacks HDR support and a high-DPI experience.

An integrated KVM switch allows you to connect to two further PCs and share a single keyboard and mouse. Further connectivity comprises a pair of HDMI 2.0 ports, plus a DisplayPort 1.4 socket. All told, it's a comprehensive package from a productivity perspective.

Given the modest 300nit panel brightness rating, the EV3895 is punchier and more vibrant than you might expect. While it's a resolutely SDR rather than HDR monitor, colors have plenty of pop and the IPS panel's inherent contrast is decent, albeit not as good as the best VA monitors in that regard.

Overall, it looks similar to other monitors based on what is almost certainly a 38-inch curved panel from LG. On the upside, that means it looks natural and accurate. The downside involves the modest pixel pitch—3,840 by 1,600 pixels is pretty generous from a pure desktop real-estate perspective. And the 111dpi pixel density means that running at 100 percent scaling should be comfortable for most people. But this clearly isn't a

high-DPI experience. If you want really crisp fonts, for instance, this isn't the monitor for you. Even still images aren't as sharp and detailed as you'd experience on, say, a 27-inch or 32-inch 4K display.

Of course, the wide aspect makes the Eizo FlexScan EV3895 great for viewing multiple documents or apps in parallel. Combine that with the KVM switch and you have real multitasking power.

As for color accuracy, it's worth noting that this isn't one of Eizo's higher-end ColorEdge models that are aimed at serious content creation pros. The 94 percent coverage of the DCI-P3 color space, for instance, is merely adequate. Likewise, the EV3895's OSD menu offers relatively limited color and calibration options, compared with those full 'pro' ColorEdge models.

All of this makes the \$1,899 asking price feel a bit punchy. There are far cheaper 38-inch options out there, even if none quite offer the same comprehensive productivity package from a connectivity perspective. —JEREMY LAIRD

VERDICT **Eizo FlexScan EV3895**

■ **WHAT A VIEW** Epic productivity feature set; really nice IPS panel.

■ **WHAT A VISTA** No HDR support; mediocre pixel density; expensive.

\$1,899, www.eizo.com

SPECIFICATIONS	
Panel size	28-inch
Panel type	IPS
Resolution	3,840 x 2,160
Brightness	300cd/m2
Contrast	1000:1
Pixel response	4ms
Color coverage	100% sRGB
Refresh rate	60Hz
Vesa	100mm x 100mm
Inputs	DisplayPort, Mini displayPort, HDMI 2.0 x2

XPG Xenia 15 KC

1440p high-refresh in a slim, slick, portable package

THE NEW XPG Xenia 15 KC gaming laptop packs some serious punch for a slim portable PC. But wait! Isn't XPG the performance brand of memory specialist Adata? Yes, it is—and the Xenia 15 KC is based on a 15-inch laptop chassis engineered by none other than Intel.

Flip the Xenia upside down and you'll find an Intel Corp sticker and product code, which is no bad thing. Intel not only has the resources to engineer a laptop right in a general sense, but it also has access to more information than anyone else when it comes to optimizing a chassis design to suit its mobile CPUs.

The chassis in question has been around in some form since 2019. It's mildly revised here and remains pretty sweet. Built largely from magnesium, it's slim, slick, well built, and relatively portable at 1.94kg. It feels solid throughout with a particularly rigid keyboard bed. Speaking of which, that keyboard sports per-key RGB lighting, complemented by a slice of RGB lighting on the front of the chassis.

Port location with the USB-C, HDMI, and LAN on the back makes sense, while three USB-A sockets, headphone jack, and the SD card slot are on the sides. Elsewhere, the specs look good, with an eight-core Intel CPU and Nvidia's RTX 3070 graphics with 8GB of RAM, plus

32GB of XPG's own RAM and a 1TB version of XPG's rapid Gammix S70 SSD. Best of all, the 15-inch IPS panel combines 1440p resolution with 165Hz refresh.

It's taken a while, but laptop makers are final getting the balance right regarding resolution and refresh. A 4K panel on a relatively thin laptop like this would be silly. 1440p is a better match with the RTX 3070 graphics while offering more image detail than 1080p. Admittedly, you might not notice much of a difference compared with 1080p in-game. But for day-to-day computing, it makes for a much, sharper, clearer display with superior font rendering than 1080p. Less impressive are the Xenia's speakers, which produce thin, muffled, and muddy sound.

The caveat to all this involves that slim chassis. It inevitably means that the RTX 3070 chip is clocked at the lower end of the available range at 1,290MHz. The compact chassis also means that fan noise at full performance is a little intrusive, which is often the case with thinner laptops with smaller, whinier fans. That said, there is a quick-jump button above the keyboard which can be used to toggle between 'Battery', 'Balanced' and 'Performance' modes. Balanced mode is much quieter than Performance while only reducing performance by 10 percent in games.

As for actual gaming performance, the Xenia puts out some pretty decent numbers at 1080p, clocking nearly 80fps in *Total War: Three Kingdoms* at Ultra quality and around 70fps at the same settings in *Metro Exodus*. Of course, this is a 1440p laptop so the best image quality will always be running native at 2,560 x 1,440 pixels. At that resolution, you're looking at frame rates in both games of about 50fps at Ultra quality.

It's also worth noting that the IPS panel looks decent interpolated to 1080p. Running online shooters at 1080p and with some of the eye candy toned down will comfortably net high frame rates, if that's more your thing.

The performance of the Intel eight-core CPU is strong. Arguably even more impressive is the 1TB SSD. It's XPG's own Gammix S70, a PCIe 4.0 drive capable of 7GB/s reads and 6GB/s writes, along with strong 4K random access performance. Add in the 32GB of XPG RAM and it's a strong overall proposition. As for battery life, the Xenia delivers over six hours in PCMark 10, which is good for this class of laptop and helps to round out what is a strong overall package. —JEREMY LAIRD

VERDICT XPG Xenia 15 KC

9 **POWERHOUSE** Nicely engineered chassis; great mix of components.

OUTHOUSE Slim proportions limit GPU power; not exactly cheap.

\$2,499, www.xpg.com

BENCHMARKS


	XPG Xenia 15 KC
3DMark Fire Strike	21,171
Cinebench R20 Multi-core (points)	5,113
GeekBench 5 (single-core)	1,572
GeekBench 5 (multi-core)	9,267
PCMark 10 (points)	7,027
PCMark 10 Battery Life (h:m)	6:25
Movie playback Battery Life (h:m)	5:40
Total War: Three Kingdoms (1080p, Ultra fps)	79
Metro Exodus (1080p, Ultra fps)	70

SPECIFICATIONS

CPU	Intel Core i7-11800H
Graphics	Nvidia GeForce RTX 3070 8GB, 1,290MHz boost
RAM	32GB DDR4
Screen	15.6-inch, 1440p, IPS, 165Hz
Storage	1TB PCIe 4.0 NVMe SSD
Ports	3x USB-A, 1x USB-C, LAN, HDMI, SD Card
Connectivity	Intel Wi-Fi 6, AX201, Gigabit Ethernet, Bluetooth 5.2
Weight	1.94kg
Size	23.4 x 35.6 x 2.1cm



XPG's Intel-engineered
15-inch chassis packs a
1440p, 165Hz punch



The AOC 24G2 doesn't break records, but it does deliver a great all-round gaming experience

AOC 24G2

AOC makes high-refresh IPS gaming affordable

GIVE IT UP FOR that rarest of things in these unusual times, a gaming peripheral that won't break the bank. We present the AOC 24G2, a 24-inch gaming monitor available for around \$180, though day-to-day pricing is pretty volatile right now

For that, you get all your basic gaming needs covered. And then some. High refresh, IPS image quality, adaptive sync, even slim-bezel styling, it's all here. At this price point, of course, something has to give. Along with the modest 24-inch proportions, you'll also have to make do with a mere 1080p resolution. But then that's a good fit with both the performance of most budget-orientated gaming PCs and the average gamer's actual budget too. More pixels make for a heavier GPU load, which is a real problem given current graphics card pricing.

At the heart of the AOC 24G2 is a 24-inch 144Hz IPS LCD panel. That bodes well for basic image quality, including colors, viewing angles, and pixel response. For pixel response, the panel is officially rated at 1ms by the MPRT metric and 4ms by the more demanding grey-to-grey measure. It's nice to see a monitor manufacturer quote both response standards. Neither is a totally reliable guide to real-world performance, but more information is always better. This is a reasonably responsive gaming panel, but it's not at the very cutting edge of IPS technology, which currently clocks in at 0.5ms MPRT and 1ms GTG.

Elsewhere, adaptive refresh takes the form of AMD FreeSync premium certification, but there's no HDR support. The latter omission isn't much of an issue, in practice. Vanishingly few monitors that claim HDR capability actually deliver and certainly none do so anywhere even close

to the AOC 24G2's price. For the record, AOC pegs the panel's color accuracy at 91 percent coverage of the Adobe RGB color space, which is adequate but nothing special, especially for an IPS monitor.

As for inputs, a pair of HDMI ports plus DisplayPort and even a VGA connector make this a widely compatible contender and a decent fit for both PC gaming as well as being a monitor for the PS5 or Xbox Series X. You also get a fairly slick, slim-bezel design and, most impressively, a fully adjustable stand that even supports rotation into portrait mode. Nice!

NOT HUGELY PUNCHY

However, if you're looking for a hugely punchy gaming panel, the AOC 24G2 isn't it. If that isn't a huge surprise given the 250 nit brightness rating, the basic image quality is actually decent. The 24G2 looks like an IPS panel for all the right reasons, with reasonably vibrant and accurate colors and decent factory calibration. There's only a little compression in darker tones visible in test images.

More importantly, it's quick. Not quite as quick as the very best 1ms IPS screens, but really not far off from a subjective perspective. AOC has included three levels of pixel-accelerating overdrive in the OSD menu and even the fastest doesn't result in any clearly visible overshoot. That means you can actually use this panel in its most aggressive mode without any obvious downsides.

The 144Hz refresh will likewise be plenty for most gamers in terms of perceived latency and input lag. You'll probably want something a bit quicker if you aspire to tournament-level esports. For everyone else, the 144Hz should keep you competitive enough.

With relatively modest brightness and lack of HDR support, not to mention the 1080p native resolution, the visual spectacle on offer isn't that memorable. The AOC 24G2 is at its best knocking out the frames in games such as *Fortnite* and *Apex Legends* rather than attempting a forensic inspection of the impact of ray-tracing in *Crysis Remastered*.

The AOC 24G2 doesn't break records or set new standards. The basic package of a 24-inch IPS panel with 144Hz refresh and 1080p resolution isn't radical. But at this price, it makes for an appealing all-purpose gaming package. —JEREMY LAIRD

VERDICT **AOC 24G2**

8 ■ ALL THAT GLISTERS Good all round feature set for the money; decent basic image quality.

■ IS NOT GOLD No HDR support; 1080p feels a bit poxy, even at this price

\$180, www.aoc.com

SPECIFICATIONS	
Panel size	24-inch
Panel tech	IPS
Native resolution	1,920 x 1,080
Response time	1ms MPRT (4ms GTG)
Refresh rate	144Hz
Contrast	1,000:1
HDR	No
Color coverage	91% Adobe RGB
Brightness	250 nits
Inputs	DisplayPort, 2x HDMI, VGA
Other	AMD FreeSync Premium



The MEG Z690 Unify is one of 13 Z690 boards from MSI

MSI MEG Z690 Unify

Seriously stealthy and stuffed with M.2 support

ANY COLOR so long as it's black. So it goes for MSI's mid-range MEG boards and the new MSI MEG Z690 Unify is no exception. While the look is familiar, the specifications have had a major overhaul to support Intel's new Alder Lake CPUs, not to mention no fewer than five M.2 sockets, dual 2.5 GbE LAN, Wi-Fi 6E, plus DDR5 and PCIe 5.0. It's quite a package.

If that's true of the feature set, it also applies to the board's physical bearing. The heatsinks have a lot of mass, as the structure reaches out over the IO area. The sheer scale of the cooling and the surface area connected via a heatpipe bode well for keeping the high-end MOSFETs running well within their specified range.

Speaking of cooling, for the CPU there are only mounting holes for the new LGA1700 socket, not the legacy LGA1200 item. The holes are similar, but the overall "Z" height is slightly lower on the newer platform. In other words, make sure to check on mounting pressure, even if you have an LGA1700 mounting kit for your cooler.

To the right of the socket are four DRAM slots capable of supporting up to 128GB of DDR5 RAM. MSI lists support up to DDR5 6666MHz using one DIMM per channel. With all slots populated, that drops to 4000MHz. Next up, there's a two-character debug LED for troubleshooting boot issues in the upper-right corner, plus a four-light EZ Debug LED for Boot, VGA, DRAM, and CPU, further helping to identify trouble spots.

The Unify sports a 20-phase VRM, with 19 phases dedicated to Vcore and an incredible 1,995A available to the processor. That's enough to support the power-hungry Intel Core i9-12900K at stock speeds and overclocked. The bottom half of the board is almost entirely covered by heatsinks and shrouds. There's a dedicated headphone amplifier (ESS-made) that handles headphones up to 600ohm, plus the usual EMI and noise mitigation measures too. As for PCIe, the top two slots source their bandwidth from the CPU, and the top slot runs up to PCIe 5.0 x16, with the second slot supporting up to x8 speeds. If the bottom slot is populated, the top drops to PCIe 5.0 x8.

Wedge above, below, and between the PCIe slots are a total of five M.2 sockets. Three support PCIe 4.0 x4, while the other one is limited to PCIe 3.0 x4. You also get V-Check Points Lite voltage read points to read the voltage through a multimeter and not software. On the rear IO area, there are no fewer than 10 USB ports, including seven USB 3.2 Gen2 ports, two USB 2.0 ports, and a USB 3.2 Gen2x2 Type-C (20 Gbps) port. Arguably, that's a bit light on the USB-C connectivity and Thunderbolt doesn't appear at all.

The BIOS menu is just like the Z590 Unify with easy access to overclocking options, though, with Alder Lake's hybrid config, the actual process has changed significantly. Of course, it's early days for the Z690 platform and Alder Lake CPUs generally, so drawing performance conclusions isn't easy.

Certainly, there were no obvious weak spots, though the Unify appears to be a power-hungry board, which isn't ideal given that high-end Alder Lake CPUs are heavy drinkers themselves. That said, VRM temperatures topped out just above 54°C during stock and overclocked testing, so there's nothing holding you back from an overclocking perspective.

All told, this is a solid board both metaphorically and physically. Oodles of M.2 storage capacity sweeten the deal and the only obvious omission is a total lack of integrated RGB illumination, though there are integrated headers on the board and MSI Mystic Light support if you want to add your own. —**JOE SHIELDS**

VERDICT **MSI MEG Z690 Unify**

- MEGA BOARD Sleek aesthetic; lots of M.2 support; great overclocking support
- MEGA BORED No RGB lighting; possibly a little power hungry

\$489, www.msi.com

SPECIFICATIONS	
Socket	LGA1700
Chipset	Z690
Form Factor	ATX
Voltage Regulator	20 Phase [19+1 105A MOSFETs for Vcore]
USB Ports	1x USB-C 3.2, 7x USB 3.2 Gen 2, 2x USB 2.0
LAN	2x 2.5 GbE
Audio	5.1, optical SPDIF
PCIe	Dual x16 5.0 (x16, x8/x8), one x4 3.0
Memory	4x DDR5 6666+(OC), 128GB Capacity
M.2	3x PCIe 4.0 x4 (up to 80mm), 1x PCIe 4.0 x4 + SATA, 1x PCIe 3.0 x4
SATA	6x 6bps (Supports RAID0/1/5/10)
Fan/Pump Headers	8x 4-Pin
Diagnostics	Two-character debug LED, four-LED Post Status Checker
Wi-Fi / Bluetooth	Intel AX201 Wi-Fi 6E, BT 5.2
Warranty	3 Years

BENCHMARKS

	MSI MEG Z690 Unify
Cinebench R23 multi-thread (Pts)	27,393
Cinebench R23 single-thread (Pts)	1,990
AIDA 64 memory bandwidth (GB/s)	79
3DMark Firestrike Extreme (Pts)	17,121
Peak system power (W)	340
Idle system power (W)	71
Overclocking P Cores (GHz)	5.1
Overclocking E Cores (GHz)	4.1

Test setup: Intel Core i9-12900K, DDR5 5200, Asus TUF RTX 3070, Windows 11 64-bit

SUBSCRIBE TODAY

Find your next great read



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

www.magazinesdirect.com/B2XMPC

Asus ROG Strix Go Core Moonlight White

A great-looking budget headset with few frills

NOWADAYS, IT SEEMS headphone features are more of a selling point than the actual sound quality. Though extras like ANC, on-the-go media controls, Bluetooth connectivity, and lift-to-pause are great, it's nice to focus our attention on what really matters.

Asus's ROG Strix Go Core gaming headset sits in the budget range of headphones, costing just under \$80. It aims to bring comfort, practicality, and great sound into one unit but, of course, at this price something has to give and so it foregoes a few frills you get on other models. So, in a world of features, how does this gaming headset stand out?

Straight out of the box, our initial impression is that these headphones certainly feel better than their \$80 price tag. The design is a crisp two-tone white and grey that Asus appropriately calls Moonlight White. The construction is a soft-touch plastic that gives it a more premium feel than expected—there are no standout RGB or colored logos here. Asus has taken a less-is-more approach that gives it a sleek overall look vibe.

Comfort is impressive too, the headset is light without feeling cheap or flimsy. It has a soft leatherette padding that cups your ears securely without too much pressure, meaning you can keep them on throughout a decent gaming session with no real issues. The same goes for the headband too, which is made of the same material and padding. There is plenty of adjustment, too, on the headband and they can also fold flat for significantly easier portability.

As we mentioned, these headphones aren't packed with features but they do have the basic controls. On the left earcup are a mute switch, volume wheel, and a microphone port. The detachable microphone is pretty impressive and can be easily removed when you aren't always on a call or in a party chat and just want to listen to some music. It's made of

sturdy unidirectional metal and comes with a foam filter tip in the box too. The pickup is pretty good and on playback, voices come through clearly.

Unlike the microphone, the 3.5mm cable is not detachable. It does come with a cable extender with headphone and microphone jack splitters, but that's perhaps less useful than a detachable cable. The 3.5mm aux cable is the only connectivity option for this headset, and so if the original cable breaks, then the headphones are useless.

FEW BELLS AND WHISTLES

As construction and overall feel go, this headset ticks a lot of boxes. It doesn't have the bells and whistles of some other models but, at this price, we wouldn't expect that. But how do these cans sound? The 40mm Asus essence drivers do a great job at pushing out a well-balanced crisp sound. The highlights are found in the high-ends and mids, it has plenty of clarity for a budget headset and certainly feels more premium than it is.

Where the sound lacks slightly is in the low-end bass. It's not flat by any means, but it could do with a little more punch at times. For instance, when playing shooters, it could do with extra bass for greater immersion in high-action elements. However, the great directional audio covers the immersion lost through the soft bass.

These cans are also pretty decent for listening to music. Due to the lightweight design, you can keep them on your head for ages and listen on-the-go with no issues. The same sound quality pros and cons from gaming also apply to music—it could do with more bass in some genres.

For its price and intended use, this headset is a well-balanced peripheral. The only minor spoiler is the aforementioned slight lack of bass, but



with an equalizer this problem can be reduced. It isn't packed with features but is compatible with devices that support 3.5mm connectivity. Comfort is great, with a lightweight build, yet it doesn't feel cheap. The design and 'Moonlight White' color scheme make it look pretty sleek, so it's equally good for using as a gaming headset and subtle enough for using out and about for music. Along with the detachable microphone, it's a great all-rounder from Asus. —SAM LEWIS

VERDICT

8

Asus ROG Strix Go Core Moonlight White

■ **MOONLIGHT DELIGHT** Crisp subtle design; ultra-light construction; clear microphone; good mid to high-end sound.

■ **DARK SIDE OF THE MOON** Slight lack of bass; fixed 3.5mm cable.

\$79.99, www.shop.asus.com

SPECIFICATIONS

Driver Type	40mm Asus Essence driver w/ neodymium magnets
Frequency Response	20Hz-40,000Hz
Impedance	32ohms
Connectivity	3.5mm auxiliary
Design Style	Circumaural, Closed back
Weight	0.54lbs

techradar.pro

IT INSIGHTS FOR BUSINESS



<http://www.techradar.com/pro>

mobilism.org



SteelSeries Apex 3 TKL

Sane or insane in the membrane?

STEELSERIES' APEX 3 TKL is a smaller take on the brand's most budget-friendly keyboard. As it's a TKL format, it takes the 10-key numeric pad away from the standard Apex 3 board, making it a favorite amongst esports competitors. With less room taken up on your desk, there's more room to swipe that mouse around. So, is this membrane keyboard a great pick-up for a price tag of under \$50? This board also has a pretty nifty feature that may separate it from its main competitors.

Although this is a TKL model, the Apex 3 TKL still offers you a taller format than normal. There is plenty of room on this board and it doesn't feel as cramped as others. From the few first clicks after setting this keyboard up, it feels a little more premium than its price tag may suggest, which is always a positive.

It's probably a little too lightweight for our liking, though, and it can slip around on your desk. Unfortunately, you may also find your fingers slipping across the keys after a really long period of typing or action-packed gaming sessions. The keycaps are made of a soft, smooth plastic that results in a lack of grip. However, for general usage, they do feel pretty decent.

The keys have a floating design that sits on an RGB panel. This way, the keyboard is pretty bright overall, which is a nice feature from a budget board. Controlled with SteelSeries Engine software, it has 8-zone RGB lighting that distributes a nice even glow. With the RGB on, it looks great, but switch it off and you can see the plastic lurking beneath, which isn't the prettiest of views. The overall design is well executed and follows the

aesthetic of the rest of the SteelSeries lineup. It features dedicated multimedia controls and a clickable volume roller on the top right of the keyboard. It has kickstands on the back that provide a comfortable writing angle although there is no wrist support, which is a shame, but not entirely unreasonable for under \$50. Underneath, you will also find some good cable management solutions with three routes for the cable to follow; left, middle and right. No one wants their setups looking like a bowl of spaghetti, so anything that can help out can only be a good thing.

So what are these keys actually like? First off, the keys feature SteelSeries whisper-quiet switches, so don't expect to find any satisfying clack sounds with this board. With this being a membrane board, they feel pretty heavy and don't give you that instant satisfaction. On the flip side, they are nice and soft for writing and it's quite relaxing to use quiet switches like these. They are good enough for casual or light competitive gameplay but don't expect to rock up to an esports tournament with this thing. They are pretty springy but the actuation isn't the shortest. That means it won't be the quickest keyboard out there, but it performs decently for the price.

As a whole, there's not too much to lose if you are looking for a budget keyboard and end up choosing this Apex 3 TKL. The board has a decent design with great RGB, a comfortable writing angle, good cable management (albeit with a fixed cable), and plenty of customization with its macro and media controls. If you have tight space for your setup, a TKL board makes the most out of your desk

space. It doesn't look like a lot, but you get more room for your mouse which is always pretty important for gaming. Membrane boards aren't always the popular choice for gaming but there are worse options for your money. The keys are soft, springy, and quiet and for typing and casual gaming.

There is one last thing, however, this board is the first water-resistant TKL gaming keyboard. If you have ever knocked a drink over your precious peripherals you'll know it's not a pleasant experience, so peace of mind is nice. As a package, the Apex 3 TKL is a good budget buy if you aren't after a clacky mechanical keyboard. —SAM LEWIS

VERDICT



SteelSeries Apex 3 TKL

■ RIDING THE WAVES IP32

water resistant; 8-zone RGB illumination; comfortable angle; multimedia controls

■ **UNDER THE WATER** Slippery and heavy keys; no wrist rest; not quite enough weight; no USB passthrough

\$44.99, www.steelseries.com

SPECIFICATIONS

Switch Type	SteelSeries Whisper-Quiet Switches
Switch Durability	20 million presses
Design Style	TKL
Connection	USB-C
Wrist-rest	N/A
Weight	1.41lbs

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

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



FUTURE



Ordering is easy. Go online at:

magazinesdirect.com



Or get it from selected supermarkets & newsagents

Asus ROG Strix Impact II Moonlight White

This mouse looks crisp in Asus's moonlight white



HERE AT *Maximum PC*, we are big fans of Asus's moonlight white aesthetic. It falls somewhere between being not too in your face but not too minimal either, so it's nice to see it appear on the ROG Strix Impact II mouse. We recently reviewed Marsback's Zephyr Pro mouse and it's fair to say that this ROG Strix Impact II is quite the opposite. It's still a mouse, just without some of the gimmicky features, such as holes in the main body.

Asus's ROG division has been around since 2006 and so it knows what it's doing with gaming peripherals. Mice are a key element of its arsenal and this Strix Impact II is the same price as the ROG Gladius II with one less programmable button. So, is this Strix Impact II a good companion for your setup and does its performance match the looks?

These days, technology is advancing so rapidly that regular items get over-engineered to a point where it just becomes gimmicky. With a mouse, there is no need to reinvent the wheel and that's where this ROG Strix Impact II comes in. On first impression, it's a clean-looking mouse with nothing too crazy going on. It features Asus Aura Sync RGB (customizable via the Armory Crate software), which is found in the wheel itself, the logo, and towards the front of the mouse. With the right color scheme, it can contrast nicely against the white and grey aesthetic of the Impact II.

It's a reasonably small mouse, not enough to put it in the mini category, but those with big hands might want to consider something a little larger. With no rubber on the sides, the Impact II has a grooved plastic texture instead as a substitute. It is fairly grippy but with some lengthy usage in heavy gaming sessions, it can get a little slippery over time. The finish is a soft matte plastic that doesn't appear to pick up too many fingerprints, but a white mouse isn't the easiest to keep clean.

The ROG Strix Impact II has five configurable buttons. Compared with the original Impact I, that means two extra thumb buttons. It's an ambidextrous mouse, yet these two extra buttons are on the left-hand side, making it slightly easier for right-handed people. It's not necessarily the best format for the MMORPG fans out there either, as it doesn't feature 101 buttons to be mapped but this mouse will have enough buttons for the general gamer. Its main right and left buttons feature Omron 50M switches and are super-light and easy to actuate with minimal pressure, but they also give a satisfying click.

It's important to note that they aren't light enough to accidentally click, but are great for gaming as it gives you a quicker response time. We all know how important this is and the shorter the input delay, the better chance you have in competitive gameplay scenarios.

REINVENTING THE WHEEL

Another solid part of this mouse is its scroll wheel. This translucent rubberized wheel feels great to the touch and provides a quiet and accurately smooth scroll. The wheel's button is strong and it can't be easily pressed. It's always frustrating when you are flying down a page and accidentally hit the scroll button to disturb your procrastination. The mouse also glides pretty well due to its lightweight nature, so you can certainly get some quick flicks out of it, which is perfect for competitive gameplay.

Below the mouse is where you will find the DPI button. Most competitive players will benefit from having this here as you can't accidentally switch DPIs mid-game. By default, it provides four different levels of sensitivity; red, pink, blue, and green, with the latter being the most sensitive. Each time the DPI button is pressed, it will briefly indicate a color to let you know the sensitivity.

The performance of the mouse feels pretty tactile and accurate. Of course, this alters depending on your sensitivity but there is lots of tuning potential in the Armory Crate software that allows you to get your specific desired results.

The 6,200 DPI sensory tracks at up to 220 inches per second. In addition to its 1000Hz polling rate, that means this mouse will provide you with a decent response and precision, but still maintain accurate control in exactly the way you would want a mouse to do.

The Impact II is priced at just under \$50, which makes it an enticing choice for many gamers. The new moonlight white color scheme looks pretty neat and seeing as it has the performance to match, it could be a good budget gaming peripheral to add to your rig. —SAM LEWIS



Asus ROG Strix Impact II Moonlight White

FULL MOON Clean design; decent lightweight build; minimal RGB; swappable switches

LUNAR ECLIPSE Not enough buttons for MMORPG players; fixed cable; plastic grips

\$49.99, www.shop.asus.com

SPECIFICATIONS

Sensor	PixArt PAW3327 Optical
Resolution	6200dpi
Speed	220 IPS
Polling Rate	1000Hz
Weight	0.17lbs (with cable)

DISCOVER THE ALTERNATIVE TO APPLE AND WINDOWS

This is a treasure-trove of Linux and open source knowledge, offering in-depth features, fun step-by-step guides, Raspberry Pi projects, security tips and much more!



FUTURE



Ordering is easy. Go online at:

magazinesdirect.com



Or get it from selected supermarkets & newsagents

AVerMedia PW513

For professional use, streaming, or both?

WE RECENTLY REVIEWED AVerMedia's PW315 webcam and were pretty impressed with the features and performance we got for a relatively affordable price tag. The Taiwanese company specializes in webcams and the AVerMedia PW513 is a more premium version of the model we've already seen.

Quality is key to good video communication. After all, there's nothing worse than your video call looking like a game of *Minecraft*. Though we all experience dodgy internet connections from time to time, one way of improving the quality is to upgrade your outdated webcam. So, having given the PW315 the thumbs-up, we have high hopes that the PW513 can improve on it.

One thing that the PW315 suffered from was its build quality but, here, AVerMedia has stepped up its game. Instantly, you get a sense that this is a more premium device—it's larger, heftier, and feels more rugged. Though you are unlikely to take it around with you often, if you do want to throw it in a bag for a Zoom meeting on-the-go, it's a durable product that's built to last.

The mount is rubberized for better grip and features a typical tripod screw hole to make this compatible for other uses. The PW513 is sturdy and the camera itself sits on a 360° swivel ball mount that not only allows you to find the perfect angle but also gives you confidence that the camera won't slip down. With a superb 94° field-of-view, you will be sure to get you and your room fully in frame.

The design is an all-black barrel look that we like a lot. We also love the privacy shutter, an arm that pivots and is super-easy to use. Having a fixed privacy shutter is much better than a removable cap, and makes for less desk clutter. It also has a cutout for the left microphone if you are just using it for audio chat. The lens sits in the middle and is a lot larger than the PW315, but we will focus (sorry!) on that later. There is a microphone on each side of the camera and also on the right side is a blue power indicator.



Round the back is the USB C 3.0 port and the camera also comes with a USB A to C 3.0 cable in the box. AVerMedia claims that the PW513 has a 'crystal clear picture quality', but the company happens to be right on the money with this bold statement.

This is a plug-and-play webcam and straight out of the box, the 4K quality is pretty impressive. Its stand-out feature is its Sony 4K CMOS sensor and, used in a well-lit room, the clarity was exactly what we would have expected. Colors pop and the fixed focus always seems to keep you detailed in the shot. The smoothness isn't perfect, as the 4K mode doesn't shoot in 60fps, but the PW513's quality overshadows the smoothness issues.

ADJUSTABLE SETTINGS

Overall, the default setup is pretty impressive, but to make matters better, AVerMedia has a dedicated program called CamEngine. This can be used to adjust the camera's settings and includes sliders for brightness, sharpness, and filters. It saves your selections onto the camera, which is perfect if you are switching between different PCs. It also shoots in 1080p 60fps for those who want more frames, the quality is still great and for most people, it will be more than good enough—though the ability to shoot 4K 60fps video would have made this an almost ideal webcam.

Another plus for the PW513 is its built-in microphones. While many will choose to pair this webcam with a headset for the best audio and microphone solution for video calling or streaming, if you find yourself without a headset, the built-in mics will certainly do.

This is a great webcam for either streaming or professional purposes—the CamEngine software even gives you the choice to optimize the setup for either usage. With a great build quality, sleek design, a wide field-of-view, and a sharp 4K sensor, this webcam is the best of both worlds. Our only reservation is that, at \$170, it is quite pricey.

However, the PW513 does exactly what you'd hope it would for that price. Maybe an HDR mode could improve the quality, but we had no issues with clarity and color pickup. With the popularity of both streaming and work video meetings, now is a great time to invest in a high-end webcam and this is certainly one to consider. —SAM LEWIS

VERDICT **AVerMedia PW513**

9 **CRYSTAL CLEAR** Great 4K image quality; wide FOV; privacy shutter and great build quality

IMPROVEMENTS NEEDED No HDR; fixed focus and no 4K @60fps

\$169.99, www.avermedia.com

SPECIFICATIONS	
Video resolution	4K @ 30 FPS, 1080p @ 60 FPS
Image resolution	8.0 megapixels
Video encoding	MP4
Field of view	94°
Focus type	Fixed
Connectivity	USB 3.0 C-A
Weight	0.86lbs

Age of Empires IV is a whistle-stop tour of historical battles



Age of Empires IV

Will this new entry in the dynasty sweep all before it?

SO MUCH TIME has passed since the last *Age of Empires* release, one that wasn't a remake or a definitive edition, the chances are that anyone reading this hasn't played a game in the series, or at least hasn't thought about it for a long time.

Things have changed a lot over the past 16 years. Series originator Ensemble Studios is no more, replaced by Relic Entertainment (*Homeworld*, *Company of Heroes*) and World's Edge, an internal Microsoft studio set up to look after *Age of Empires*. Things have also changed outside of the development studio, especially in the world of *Total War*. The two series overlap in more than just timeframe, with Creative Assembly's behemoth pushing the boundaries in terms of units on-screen and, more recently, fire and magical effects.

You won't find any of the latter in *Age of Empires IV*, unfortunately, but then that's not the nature of the game. Presented as a historical documentary, with live-action cut-scenes full of foreboding ravens fluttering around on battlefields and diagrams of how troops clashed in real wars, *Age of Empires IV* takes itself pretty seriously.

It all begins at Hastings in England, where a battle was fought in 1066 that

changed the course of a continent. Duke William of Normandy's victory over King Harold of England is a foregone conclusion, so the introductory campaign sees you play as William, consolidating his power and mopping up resistance.

This is followed by three campaigns set in less clear-cut times: the Hundred Years War, 116 years of battles between the English and French in which both sides won significant victories; the expansion of the Mongol Empire across Asia and into Europe; and the rise of Moscow to dominate what we now know as Russia.

THINGS THAT DIDN'T HAPPEN


There's plenty of scope for the series' trademark build-up of forces, though an early rush with specialist units is still a valid tactic, the whole thing modeled after *Age of Empires II* more than any other entry. Outside of the campaign, there are eight lightly asymmetrical civilizations to be thrown at one another in skirmish maps and online multi-player.

These are naturally where you'll head for historical wish fulfillment, with the campaign rather tied to the record of how things really happened. Having the Delhi Sultanate stomp all over the English in battles that never happened makes you


wonder if a little more imagination could have been applied to the campaign.

No matter. The game may be extremely serious (though we make no apologies for finding King William's battle cry of "Incontinent!" hilarious) but it's also polished and solid. Loading times can be a little long, even from an SSD, but what you get after the wait, especially in 4K, is worth it. There's just enough realism, units are easily picked out, and some nice touches such as the sheep who obediently follow your scouts to their doom and the softly glowing crown that floats above your monarch character. Even when you have a thriving town full of villagers hunting and gathering, with military buildings pumping out units and others carrying out research, it's easy to stay on top of things just by looking at what's there.


New features aren't exactly thick on the ground, though there's now a *Homeworld*-like ability to change the shape of an assembled unit from a line to a wedge to a loose scattering. Building multiples of the same building is now the way to access unique units, such as the powerful English longbowmen who require three archery butts, all selected at once, to begin building.




No slacking here, but the idle villager button is patchy




Change your unit from a line to a wedge or loose scattering




Build settlements with farms, houses, and blacksmiths



Here you can fulfill historical wishes, however fanciful



Choose battles from 1066 to the rise of the Mongol Empire



Enlist villagers for a cavalry charge, or just go it alone

Age of Empires was just fine the way it was, and the acceptance of this on the part of the developers means that this latest game is easy to pick up and play, rather than something that drowns you in micro-management or strategy.

WORK FOR IDLE HANDS

So while roads appear automatically between areas well-traveled by villagers or troops, they don't confer a movement bonus. You can pack your buildings in close together, or put farms between blacksmiths and houses, without hitting that awful RTS state where the building's footprint turns red and won't go where you're certain it should fit.

Villagers have to be managed to prevent them from becoming idle (the idle villager button returns, but occasionally dumps the camera in the middle of nowhere), though sometimes are overcome by a fit

of usefulness and begin harvesting wood after building a lumber camp.

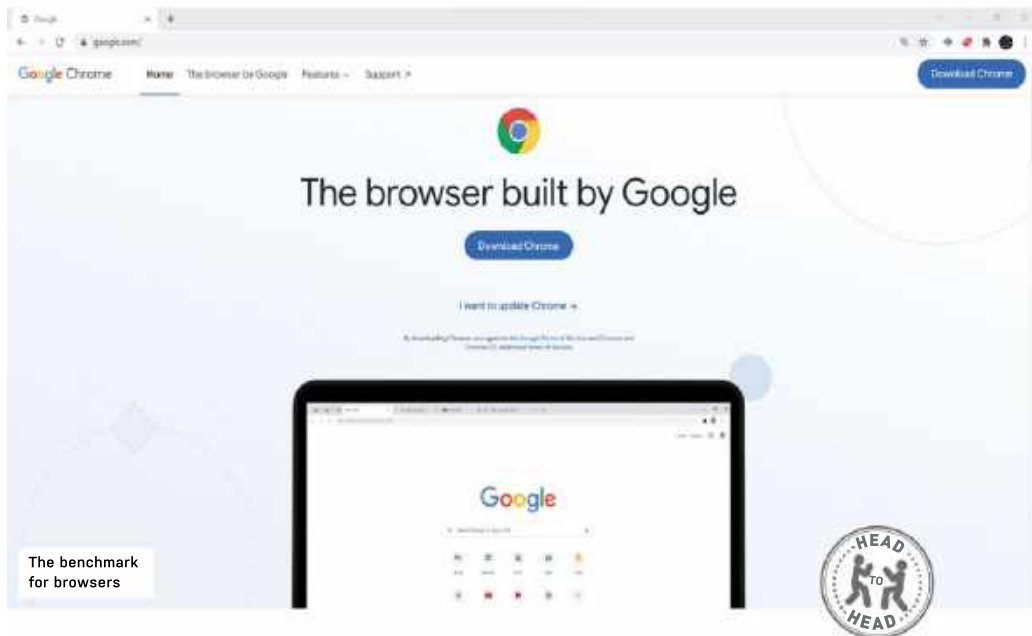
There are, of course, things that could be further streamlined, such as having to individually click on every treasure chest you uncover, or the selection boxes that pick up villagers with your cavalry just before a charge. Military units have only attack and stop commands, and it's possible for scattered archers to spend more time burning enemy fields than actually shooting arrows where they'd do the most good. If you're going to prosper in a skirmish, then learning your faction is essential—some get skills for free that others must research, unique units abound, and auto-harvesting buildings, even movable town centers, bring a whole new layer to the game.

On the whole, though, it's great to see *Age of Empires* back. Whether you have memories of sitting up all night with *The*

Age of Kings or have come to the series with *Dawn of War* and in anticipation of *Company of Heroes 3*, there's a satisfying richness about the depiction of this time period and the application of the game's particular way of doing things.

It might not make this a classic game but will certainly delight anyone looking for that medieval RTS fix. —IAN EVENDEN

VERDICT 8	Age of Empires IV
	CONQUERER Polished and solid gameplay; nicely detailed graphics; less micro-management
	WAVERER Campaign tied closely to history; villagers need too much managing
	RECOMMENDED SPECS CPU, i5 six-core / R5 1600; RAM, 16GB; GPU, GeForce GTX 970 / Radeon RX 570
\$59.99, www.ageofempires.com, T-rated	



Google Chrome vs Brave

The battle of the desktop Chromium browsers

WHEN GETTING A NEW system or finishing a new PC build, one of the first tasks is usually to change the default web browser. Yes, Microsoft provides you with its Edge browser. But for most people, it simply doesn't cut it. Beyond the simple task of downloading Google Chrome, that tends to be pretty much it for Edge. In the majority of cases, you will probably find yourself switching to the Google Chrome browser. While there are many alternatives out there, including Firefox, Opera, and Safari (for those using macOS), Chrome remains the most popular browser.

Google Chrome is built on the same Chromium open-source browser used by many other browsers, including Microsoft Edge. This means some of the features that make Chrome so popular can also be seen on other browsers. One such browser is the new(ish) kid on the block, Brave. This puts the emphasis on privacy features and has a few other tricks up its sleeve too. So, how does this new Chromium browser measure up to Google Chrome?

Browsers are like banks in that most people choose one and then don't tend to switch over from one to the other. In a way,


it's no surprise they inspire some degree of loyalty—we spend much of our lives online and browsers are one of the most accessed pieces of software. When you become used to one particular interface and its overall feel, changing browsers can be daunting. Part of the attraction of the Chrome browser is its overall feel and design. It allows for plenty of customization through its appearance settings and themes from the Chrome web store. Accent colors and background images can easily be changed and in just a few clicks, you can get a look that is tweaked especially for you.

But with Brave being built on the same Chromium platform, the same applies to this browser too. As clean installs, both layouts are similar—Google's aesthetic over the past decade has somewhat followed Apple, with a sleek, minimalistic design making the interface less cluttered and easier to navigate. This contrasts with browsers of the mid-2000s, such as Internet Explorer, which were strewn with bookmarks and extensions. These extra features are now hidden away under dropdown menus that can easily be hidden for a cleaner experience.

At first install, the default look of the Brave browser just pips Google Chrome and takes the win for the overall design. Both have a similar layout, with tabs starting at the top left of the screen. Underneath that is the search bar and to the left are back, forward, and reload buttons. On the far right, you will find the extensions and settings on both browsers. The difference here is that Brave's search engine is centered and features two unique options that we will touch on later.

Brave's landing page gives you the choice to add widgets and a more personalized feel overall. You can add a clock, news headlines, browsing stats, cryptocurrency flashcards, and more. This looks similar in design to Windows 11's widgets panel and flows down the center on frosted panels, but it manages to squeeze all this in without looking bloated.

Chrome's landing page is a Google search engine bar with a list of shortcut icons. It's a timeless look, but Brave's take is both more exciting and modern looking, whilst maintaining a clean minimal feel. Brave also uses Chrome's web store so there are plenty of customization options.



Features are one of Brave's strongest points. It takes the formula of Chrome and adds to it, providing the user with a lighter browser that feels speedier and more focused. We mentioned earlier how many people tend to stick with their browser and, in part, that's because making the transition to a new one can be difficult. But with both browsers underpinned by Chromium, switching between Chrome and Brave is a breeze. Brave supports all Chrome bookmarks and extensions. So, when you download it, you can import your previous settings, including history, bookmarks/favorites, saved passwords, and extensions. Switching is quick and easy, and you will soon feel at home again.

BRAVE NEW WORLD

In Brave, some of the most important extensions to improve Chrome, such as ad-blockers, are already pre-installed. Called Brave Shield, you will find its lion icon to the right of Brave's search bar. Brave Shield has plenty of options to tweak the ad-blocker. You can change the severity to aggressive, standard, or turn it off altogether. You can block scripts, cookies, and fingerprinting and see how many trackers have been blocked. This data can also be viewed on Brave's landing page. We spend much of our lives on the internet, so keeping your data secure is vital but, unfortunately, Chrome doesn't offer this added security feature as default.

Another standout feature of Brave is its crypto reward scheme, which aims to give users more control over their adverts and actually pays you to keep them turned on. Brave's Private Ads system rewards users through its crypto-asset BAT (Basic Attention Token). The best thing about this is that it's completely optional, like the ad-blocker, they are just there if you wish to use them. You can continue to earn more BAT or gift it to content publishers you like, similar to any other crypto asset. Brave also features a built-in wallet, called Brave Wallet, which offers you a secure way to store your BAT or other crypto assets.

Unlike Chrome, Brave applies the modern touches to a browser whilst keeping the user secure and in control. For simply searching the web, Chrome is ideal—it's a dependable browser that has been more than enough for many users over the years. However, if you like the way Chrome feels and performs but care about where your data is going and how it is being used, Brave is the better choice. The tracker on the landing page shows you how many trackers and ads have been blocked as well as how much time you have saved without using ads.

Brave offers a similar experience to Chrome but has more features and loads pages times three to six times faster. So, not only do you get a smoother browsing experience, but you also get peace of mind with an ad-blocker and pop-ups disabled.

Chrome has paved the way for modern web browsers and has proved a popular choice for a long time. However, in times of heightened online security fears, maybe it's time for a new browser that offers a better and quicker package? —SAM LEWIS

VERDICT **Google Chrome**

8 **SHINY** Reliable; sets standards for browsers; quick and tidy

SCUFFED Bland from default; no built in ad-blocking; lacking proper security features

OPERATING SYSTEM Windows 7, Windows 8, Windows 8.1, Windows 10 or later. CPU, Intel Pentium 4 processor or later that's SSE3 capable

\$Free, www.google.com

VERDICT **Brave**

9 **BRAVE NEW APPROACH** Dependable Chromium platform; new modern look; feature packed; great security; speedy browser

LITTLE SHY Not quite original, still nowhere near as large as Chrome...

OPERATING SYSTEM Windows 7, Windows 8, Windows 8.1, Windows 10 or later. CPU, Intel Pentium 4 processor or later that's SSE2 capable

\$Free, www.brave.com

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

- > 8K gaming
- > HDD performance issues
- > Integrated graphics

8K gaming...?

What would be the best and cheapest build for an 8K 60Hz HDMI Windows 11 gaming desktop?

—Shane Johnson

MAXIMUM PC EDITOR, JEREMY LAIRD RESPONDS: The first question you need to ask yourself, Shane, concerns the notion of 8K gaming. Are you sure that's what you really want? Sure, 8K sounds sexy, but it makes for no fewer than 33 million pixels, each of which has to be rendered, rasterized, and kicked out to the display. The net result is a truly debilitating workload for any GPU, even the latest and greatest.

If you are intent on the 8K prize, the best starting point is an 8K TV. Full 8K monitors arguably don't make much sense for gaming. Dell's 8K monitor, for instance, is a 32-incher and therefore arguably too small to really do justice to 8K content. Just imagine a 16-inch 4K gaming monitor. Exactly!

What you want is an 8K TV. Go for a 55-inch model—search online and you'll have the choice of LG and Samsung sets

for around \$1,000. Not exactly cheap, but far more affordable than the \$4,000 Dell. Intriguingly, a 55-inch 8K TV has the same pixel density as a 27-inch 4K monitor, so it's essentially four 4K 27-inch panels fused together.

Anyway, you're also going to need a graphics card with HDMI 2.1 support to drive 8K and 60Hz. That means either Nvidia's RTX 30 series or AMD's RX 6000 series. Either way, it will have to be one of the higher-end models which currently cost another \$1,000 minimum.

Also bear in mind that, as far as we know, all currently available 8K TVs have BGR rather than RGB subpixels. In-game, that probably doesn't matter, but on the Windows desktop, it makes for fuzzy fonts among other image quality issues. All told, 8K gaming remains at the early adopter stage. It's tempting, and if it weren't for current GPU supply issues, it would actually be reasonably cost-effective. But as things stand, the combination of cost and technical niggles make a convincing case against it.

Integrated graphics

Thank you for your recent article comparing AMD and Intel integrated graphics. For years, I have been waiting for that *Maximum PC* article that says that if I get these certain components, I can build a rig that doesn't need a graphics card. Apparently, the wait continues, though huge gains have been made.

A few years ago, I bought *No Man's Sky*. The first time I started it, I got an error message that my graphics card only had 1GB of VRAM and needed more. I've seen games that throttle down settings for low-powered rigs, but never one that wouldn't work at all.

I eventually installed a graphics card with 6GB of VRAM and the game worked fine. Are there VRAM issues with integrated graphics?

—M. Hamm

MAXIMUM PC EDITOR, JEREMY LAIRD RESPONDS: It's true that integrated graphics is improving fast, but there's still a long way to go before it gives a good gaming experience. Much of that is down to the memory subsystem, but not just VRAM quantity.

Existing PC integrated graphics rely upon a CPU-spec controller to access memory. That's bad for several reasons. First, CPU memory controllers offer less bandwidth, both because they are narrower (typically 128-bit versus up to 512-bit for conventional GPU controllers and far more for GPUs with HBM memory). Also because CPU DRAM operates at lower frequencies than GPU VRAM.

Making matters worse, that limited bandwidth (by GPU standards) must be shared between the CPU and GPU. For example, Intel's latest uber CPU, the Core i9-12900K, tops out at around 75GB/s of bandwidth using the latest DDR5 memory. Meanwhile, Nvidia's GeForce RTX 3090 GPU enjoys 936.2GB/s of bandwidth. There's no comparison.

With a shared system, it's not just bandwidth that's under contention but capacity too. The net result can be inadequate memory available to the GPU, though adding RAM can be a straightforward solution.

Of course, it's not quite as simple as fitting out

∨ submit your questions to: editor@maximumpc.com

a GPU-spec controller to a CPU with integrated graphics. The GDDR memory used with GPUs has great throughput or bandwidth but poor latency, while the DDR memory that's hitched up to CPUs is the opposite—low latency but limited throughput. The solution would be a best-of-both-worlds combination of the bandwidth of VRAM with the low latency of DRAM.

Actually, it's more or less been done. Apple's new M1 Max SoC has both CPU cores and GPUs that share a single memory subsystem offering 400GB/s of peak throughput, courtesy of LPDDR5 RAM running over a 512-bit controller. It's hard to compare graphics performance between Mac and PC, but the integrated GPU in the M1 Max is as fast as an Nvidia RTX 3060 desktop GPU and so far faster than any existing PC integrated graphics.

Another example is the latest games consoles, which have integrated GPUs sharing memory bandwidth with CPU cores. However, both the Microsoft Xbox Series X and the Sony PS5 use GDDR6 graphics memory shared between the CPU and GPU. That works for a games console but could be problematic for a PC with add-in memory.

In a console, the memory can be packaged close to the SoC, helping reduce latency. All of this means that high-performance integrated graphics is possible on the PC, but will probably require some loss of configurability.

Achieving performance levels good enough for high-quality gaming will likely require CPU and system memory to be packaged and sold together as a unit, preventing users from upgrading one but not the other.

Hard drives and the theory of speed

I am trying to transfer pictures from one HDD to the other, but I am having problems with the transfer. They are supposed to be able to transfer at the rate of 6GB/s, but it is taking hours. I wonder why these drives never seem to run at that speed? My system includes an MSI X470 AMD Gaming M7 AC mobo, AMD Ryzen 7 2700X CPU, Corsair Dominator 32GB 3000MHz RAM, 2x Seagate 500GB SSHD ST500LM000-1EJ162, all running Windows 10 64-bit.

—Edward

MAXIMUM PC EDITOR, JEREMY LAIRD RESPONDS: Part of the problem here is theoretical versus actual real-world speed. Your drives are likely to be based on the SATA III protocol with a theoretical maximum throughput of 6Gb/s—that's Gb/s, not GB/s, which are an order of magnitude apart.

Anyway, that 6Gb/s refers to the theoretical maximum speed of the SATA III protocol rather than the hard disk itself. It determines the maximum bandwidth available to the drive. The actual bandwidth available is a little lower due to protocol overheads. In theory, SATA III provides around 600MB/s of bandwidth, but the fastest transfer speed is around 550MB/s to 560MB/s. Of course, those kinds of figures are achieved by solid-state drives. Typical HDD speeds are 200MB/s or lower for both reads and writes. Even the fastest HDD yet, the Seagate Mach.2, designed for data centers, tops out at 524MB/s, but only for data stored at the edge of the platters.

Of course, even around 150MB/s should be enough for transferring a lot of images fairly quickly. Slightly confusing the, er, picture here is that

your drives are 500GB Seagate SSHDs or hybrid drives with 8GB of MLC NAND memory designed to improve performance to solid-state levels in certain scenarios. That NAND memory will help with short, sharp random disk access, the kind of traffic that dominates day-to-day computing. However, for sustained transfers, the underlying performance of the hard drive is exposed.

That said, Seagate claims your drives are good for average data throughput of 100MB/s, which ought to be enough for transferring files at a reasonable pace. The fact that it's taking hours implies either a hardware issue or some kind of configuration problem.

If your SSHD drives are old, the NAND memory may be wearing out. Another candidate could be your SATA configuration. It's worth checking that the drives are attached to the primary chipset controller rather than an add-in controller if your motherboard has multiple SATA controllers. Refer to your motherboard manual for the optimal SATA configuration for multiple drives. Drive contention can impact performance dramatically.

Also, make sure that your motherboard's BIOS is running in the correct mode for your drives. Typically, the default configuration will be correct, so resetting to global BIOS defaults is the easiest option. However, be sure to get a snapshot of your BIOS settings before making any changes at all. Depending on whether you built and set up your system yourself, there may be settings you aren't aware of that are required for a successful boot.

If in any doubt, the safest option is not to touch the BIOS settings and get some expert help instead.

Baffling blueprints...

On the basis of *Maximum PC* content generally being several notches above my pay grade and there being no fool like an old fool, please consider this when answering my question. In "The Builds" at the back of the magazine, why does the Intel Ingredients list always show GPU: AMD Radeon for Mid Range and Turbo builds? Yet the AMD list shows Nvidia? I do learn an immense amount from the magazine, but must be missing something.

—J. Watson

MAXIMUM PC EDITOR, JEREMY LAIRD RESPONDS: Not at all. There's a lot to parse in the Blueprints section. The most important aspect to grasp is that the 'Intel' and 'AMD' builds refer to the vendor responsible for the platform, rather than GPU. In practice, that means an Intel or AMD motherboard chipset and CPU. Into that, any compatible and relevant GPU can be slotted. Next year, that will hopefully also include Intel's new graphics cards, for which we have high, though cautious, hopes.

These builds are constantly under review, which is a tricky task of late given the difficulty and expense of sourcing graphics cards. But at any given moment and price point there tends to be a GPU that stands out as delivering the best combination of value and performance. At the high-end, for instance, Nvidia's RTX 3090 boards are arguably overpriced, so the AMD Radeon RX 6900 XT gets the nod.

There's no intention to have exclusively AMD or Nvidia GPUs in any given slot, but graphics card refreshes don't come along often, so changes in that area will necessarily be periodic rather than routine! 🕒

THE BUILDS

THIS MONTH'S STREET PRICES...

BUDGET



INTEL'S 12TH-GEN ALDER LAKE CHIPS are finally here, but they don't make it into the budget build. On top of this disappointing news, our i3-10100 is currently out of stock, so we made a slight upgrade to the Core i3-10105 with Intel's UHD Graphics 630. It's a similar chip, so it's not a big change to the overall build, although the i3-10105

performs marginally better with a slightly higher clock and turbo speed. At the time of writing, it has a respectable \$40 saving to its normal \$170 price tag, so it's a decent time to pick one up.

Three parts find themselves on both the Intel and AMD budget systems. Due to a stock change, we had to change our memory around on the AMD budget build and the Intel budget build. For just \$10 more, we swapped out the Crucial Ballistix CL16 for GSkills Trident Z Neo CL 16 RAM. This is the same spec as the Crucial Ballistix RAM, so was the most suitable we could find at the price.

The next change was to the PSUs. For the Intel build it was previously the 500W Aresgame AGV500 at an 80+ Bronze efficiency. On team AMD, the 500W Evga BA was the choice with an 80+ Bronze efficiency rating. This time, these have both been swapped around for the ever-so-slightly more powerful 550W Corsair CV550 unit that has the same 80+ Bronze rating.

Both builds also got an upgrade in HDDs from 1TB to 2TB for a whopping increase of \$8. With a sale on Newegg bringing these down from \$70 to \$47 it was near impossible not to make the change, especially considering the 1TB variant is sitting at \$39. SSDs remained the same as they still perform strongly in the rigs and the prices haven't shifted, which is always good to see.

The last major change for these machines is the availability of GPUs. It is hard to recommend graphics cards that cost nearly double the MSRP, but they are there for those who don't want to wait for more stock to come out. It's annoying for everyone, but at least there is some street price stock, rather than no stock at all.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Corsair 4000D Airflow	\$95	
PSU	5550W Corsair CV550 80+ Bronze NEW	\$54	
Mobo	ASRock B550 Phantom Gaming 4/ac AM4	\$100	
CPU	AMD Ryzen 5 3600	\$240	
GPU	ASUS TUF Gaming Radeon RX 6700 XT OC 12GB	NEW \$480	\$1,092
RAM	16GB (2x8GB) G.Skill Trident Z Neo CL 16 @ 3600MHz	NEW \$95	
SSD	512GB XPG GAMMIX S50 Lite PCIe 4.0 M.2 SSD	\$70	
HDD	2TB WD Blue 7200 HDD	NEW \$47	
OS	Windows 10 Home 64-bit OEM	\$32	

Approximate Price: \$1,213 or \$1,825

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Corsair 4000D Airflow	\$95	
PSU	5550W Corsair CV550 80+ Bronze NEW	\$54	
Mobo	MSI B460M PRO-VDH WIFI Micro ATX LGA1200	\$109	
CPU	Intel Core i3-10105 NEW	\$130	
GPU	MSI Ventus GeForce RTX 3060 12GB NEW	\$490	\$810
RAM	16GB (2x8GB) G.Skill Trident Z Neo CL 16 @ 3600MHz	NEW \$95	
SSD	500GB Samsung 980 NVME M.2 PCIe 3.0 SSD	\$60	
HDD	2TB WD Blue 7200 HDD	NEW \$47	
OS	Windows 10 Home 64-bit OEM	\$32	

Approximate Price: \$1,112 or \$1,432


 MID-RANGE

WE HAVE SOME BIG CHANGES for the mid-range builds this issue. We have GPUs for both Intel and AMD, however, keep in mind that these are at a street value price. For team Intel, we can bring in the Gigabyte Aorus Radeon RX 6800 XT card, which is plenty of power for this rig. We also have something exciting to pair with this GPU.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Lian Li PC-O11-Dynamic	\$150	
PSU	750W EVGA SuperNOVA GA 80+ Gold NEW	\$94	
Mobo	Asus AM4 TUF Gaming X570 -Plus WiFi	\$195	
CPU	AMD Ryzen 5 5600X	\$289	
Cooler	280mm Cooler Master MasterLiquid ML280 RGB	\$55	
GPU	MSI Ventus GeForce RTX 3070 8GB GDDR6 NEW	\$780	\$1,040
RAM	16GB (2x 8GB) Team Vulcan Z CL18 @ 3600MHz	\$63	
SSD	500GB Samsung 980 Pro NVME M.2 PCIe 4.0 SSD	\$110	
HDD	2TB WD Blue 7200 HDD	\$47	
OS	Windows 10 Home 64-bit OEM	\$32	

Approximate Price: \$1,815 or \$2,075

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Lian Li PC-O11-Dynamic	\$150	
PSU	750W EVGA SuperNOVA GA 80+ Gold NEW	\$94	
Mobo	MSI PRO Z690-A LGA1700 NEW	\$209	
CPU	Intel Core i5-12600K NEW	\$300	
Cooler	Enermax Liqmax III 240 RGB	\$80	
GPU	Gigabyte Aorus Radeon RX 6800 XT Master 16GB GDDR6 NEW	\$899	\$1,500
RAM	16GB (2x 8GB) Corsair Vengeance RGB Pro SL RGB CL18 @ 3600MHz	\$99	
SSD	500GB Corsair Force MP600 M.2 PCIe 4.0	\$90	
HDD	2TB WD Blue 7200 HDD	\$47	
OS	Windows 10 Home 64-bit OEM	\$32	

Approximate Price: \$2,000 or \$2,601

Intel's new Alder Lake 12th-gen chips are finally in the builds, specifically the i5-12600K which is a pretty sweet CPU. It comes with 10 cores combined with Intel's new hybrid architecture. Split into performance cores (6) and efficient cores (4), foreground tasks are allocated to the performance cores and background tasks are set to the efficient cores. The Alder Lake chips also integrate a new one-click overclocking feature to the BIOS that will help users get more out of their systems. The maximum turbo frequency of this new chip is 4.90GHz too, which is pretty decent considering the 12600k is around \$50 cheaper than the previous Intel Core i7-11700K.

Talking of the BIOS, the motherboard has to change to support the new CPU. MSI's PRO Z690 board arrives with its fancy new LGA1700 socket. It's a decent board but quite pricey. At an extra \$80, it's the best we could do with the new chips. Other than these changes and a change of PSU to the EVGA SuperNOVA GA, the rest of the Intel build remains the same.

Things aren't quite as exciting for the AMD system, apart from the graphics card. MSI's 8GB Ventus GeForce RTX 3070 finds itself in the AMD mid-range build this issue, again at street price.

The only other change is the same as on the Intel build, out with Corsair's RM (2019) PSU and in with the 750W EVGA SuperNOVA. They are both 80+ Gold rated but, due to stock issues, the EVGA model slots in. It is 750W this time instead of 650W from the Corsair offering, so this helps a little with future-proofing. There is good news about AMD's Ryzen 5 5600X, as this is down another \$20 from the last issue at \$290 on Newegg. And we can't resist a good sale.



TURBO

IN A SIMILAR FASHION to the Intel mid-range build, we are bringing in the top-of-the-range Alder Lake offering. Yes, that's right, the Intel Core i9-12900K—technically the world's fastest gaming chip. Just like the i5-12600K, this CPU features the new hybrid architecture, with eight performance and eight efficiency cores for a straight

50/50 split. Compared with Intel's previous top offering, the Intel Core i9-11900K, the new chip has double the core count and eight more threads. With claims that it is up to 50 percent faster, even when running on DDR4, we had to put this in the build. There is one catch though, the price. Retailing between \$590 and \$600, street prices already start from \$670.

Like the mid-range build, we had to introduce a Z690 board, but with some tasty new DDR5 RAM to bump the price up. Corsair's Vengeance sticks come in at \$300 but, hopefully, this will drop soon. For the mobo, we went for the MSI PRO Z690-A WiFi LGA1700, an efficient board to support the new CPU and DDR5 memory.

On the AMD side, we can't yet justify the jump to the Ryzen 9 5950X as that's around \$720 but have upgraded to the Ryzen 9 5900X. At \$484 on sale at Newegg, it'd be a shame not to grab this Ryzen 9 chip. Both machines come with Corsair's iCUE H150i Elite Capellix AIO. This 360mm cooler will do a great job of keeping both of these high-performance chips nice and cool. Also on both setups is the same 6TB WD Blue 5400 HDDs. With a sale still on, this is a great way to get a huge capacity secondary storage on both systems. One last similarity is that the same PSUs from the mid-range have made it into the turbo builds. At under \$100, the 750W PSU offering from EVGA is hard to ignore with its 80+ Gold efficiency.

Finally, for the graphics cards, the same situation applies to the turbo-build GPUs as with the others. They are available for a hefty premium and if you can afford to pay the street prices, you will certainly get great performance. Knowing what these cards "should" retail for, that can be a hard pill to swallow. But what alternative do you have, right now?

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	\$160	
PSU	750W EVGA SuperNOVA GA 80+ Gold	NEW \$95	
Mobo	Gigabyte X570 GAMING X	\$188	
CPU	AMD Ryzen 7 5900X	NEW \$484	
Cooler	Corsair iCUE H150i ELITE CAPELLIX AIO 360mm	\$159	
GPU	Gigabyte GeForce RTX 3080 Gaming OC 10GB	NEW \$700	\$1,500
RAM	32GB (2x 16GB) G.Skill Ripjaws V 32GB @ 4000MHz	\$135	
SSD	1TB Samsung 980 Pro M.2 PCIe 4.0	\$185	
HDD	6TB WD Blue 5400 HDD	\$110	
OS	Windows 10 Home 64-bit OEM	\$32	

Approximate Price: \$2,248 or \$3,048

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	\$160	
PSU	750W EVGA SuperNOVA GA 80+ Gold	NEW \$94	
Mobo	MSI PRO Z690-A WiFi LGA1700	NEW \$240	
CPU	Intel Core i9-12900K	NEW \$590	\$670
Cooler	Corsair iCUE H150i ELITE CAPELLIX AIO 360mm	NEW \$159	
GPU	ASRock OC Formula Radeon RX 6900 XT GDDR6 16GB	NEW \$999	\$1,550
RAM	32GB (2x 16GB) Corsair Vengeance DDR5 @ 5200MHz	NEW \$300	
SSD	1TB Gigabyte Aorus Gen4 7000s M.2 PCIe 4.0	\$190	
HDD	6TB WD Blue 5400 HDD	\$110	
OS	Windows 10 Home 64-bit OEM	\$32	

Approximate Price: \$3,174 or \$3,805

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US, Inc., 11 West 42nd Street, 15th Floor, New York, NY 10036, USA. Website: www.futureus.com. Periodicals postage paid at New York, NY, and at additional mailing offices. Newsstand distribution is handled by Curtis Circulation Company. Basic subscription rates: one year (13 issues) US: \$30;

Canada: US\$30; Foreign: US\$50. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R129220688). PMA #40612498. Subscriptions do not include newsstand specials. POSTMASTER: Send changes of address to Maximum PC, PO Box 2024, Langhorne, PA 19047, USA. Future US, Inc. also publishes MacLife, and PC Gamer. Entire contents copyright 2021, Future US, Inc. All rights reserved.

Reproduction in whole or in part is prohibited. Future US, Inc. is not affiliated with the companies or products covered in Maximum PC. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of Maximum PC. Products named in the pages of Maximum PC are trademarks of their respective companies. PRODUCED IN THE UNITED STATES OF AMERICA.



PAPER POWER

65% of the energy demand at U.S. pulp, paper and paper-based packaging mills is met using renewable, carbon-neutral biomass energy.

Discover the story of paper
www.lovepaperna.org

Source: American Forest & Paper Association, 2018



Love Paper is a registered trademark for Two Sides Ltd. Registered in the U.S., UK and other countries and used with permission.



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.

