MAXIMUMPC

MINIMUM BS • OCTOBER 2021 • www.maximumpc.com

WINDOWS 11

Microsoft's new OS finally hits our lab

- ☑ **Complete step-by-step install guide**
- ☑ **All-new super slick user interface**
- ☑ **Security nightmare broken down**

G.SKILL

FUTURE

Digital
Edition

table of contents | where we put stuff

inside
OCTOBER 2021

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.
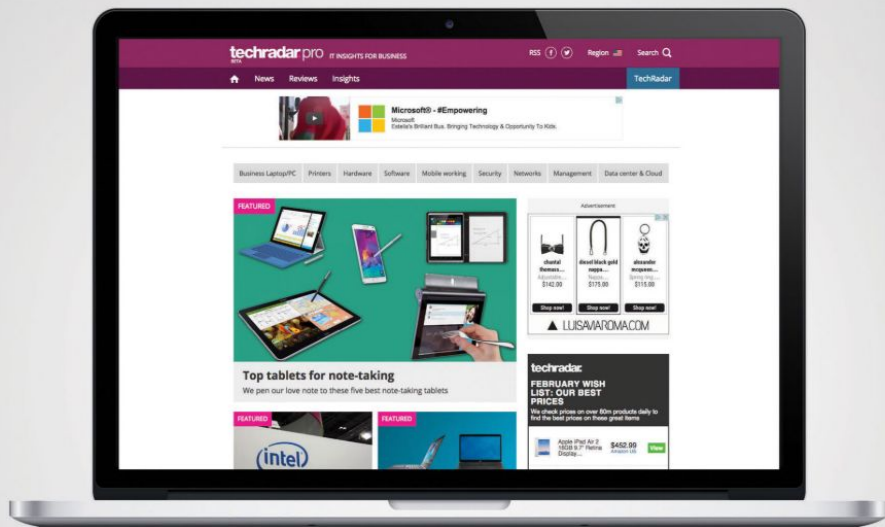
Discover the story of paper
www.lovepaperna.org

LOVE PAPER

MAXIMUMPC

Zak Storey

A NEW ERA… AGAIN

REMEMBER WHEN Microsoft said that Windows 10 was going to be the last version of Windows you'd ever need? Yeah, I remember that. Good times, September 30th, 2014. I was there, watching the announcement, just some nerdy little kid, without a job, fresh out of college. I was so excited, I'd tried Windows 8 and 8.1 Home for some years, and put up with the awful tiles and lack of a start menu, before converting the thing to Start8 and basically turning it into exactly the same as Windows 7, but 10 seemed exciting. I immediately went out and bought myself a version of Windows 8.1 Pro (with a student discount), and eagerly awaited the launch of the new operating system.

And it came, and it was great. It fixed all the flaws of Windows 8 and then some. Long gone was the ambitious "tablets are the future" OS. Finally, a practical return to form for the prestigious company.

How times have changed. When Windows 11 was first announced, it's fair to say the tech press were fairly taken aback. There's a lot to cover, a ton of new features in it, including stuff to amp up your PCs, and make them more secure, a super clean new design, and, of course, some seriously deluded security decisions (you can read more on my thoughts on that on page 11). So, I set Sam to task to give us the absolute low-down on this new plucky underdog from Microsoft, and to give you beautiful people a full-on guide on how to get Windows 11 running on your own machines as well. Is it all it's cracked up to be? Or just a MacOS-looking reskin of a once prestigious OS environment?

Windows aside, in the last issue, I promised a deep-dive feature into the connectivity standards of tomorrow,

it was something we briefly looked at in Jarred's feature on Intel's next processors, and something I wish we had spent more time on. So for this issue, Christian set out to find out everything he could on PCIe 5.0, DDR5, USB4, and everything else in between, for a stats-packed, eight-page goodie.

And for the more historically minded folk out there (or the doomsday preppers, which might also include me in their ranks after reading this), Ian went out and did his very best to investigate the Carrington Event back in 1859, the worst recorded solar storm in human history.

It was a moment in time that fried electrical cables, set telegram machines on fire globally, caused aurora to reach as far south as the Caribbean, and makes the Quebec blackout look like a pebble in comparison. In our hyper-interconnected world of today, what would a solar storm look like? What would it do, and are we prepared for it? With NASA and ESA on board, it's one to check out.

Of course, we've got even more than all of that goodness, including a barrage of tutorials, reviews, columns, news, and more, from some of the best tech writers in the biz!

I do hope you enjoy this issue, it's been a heck of a fun one. Until next time!

Zak is Maximum PC's *editor-in-chief and long-time staff member. He's been building PCs since he was 10, and is more than capable of butting heads with the biggest names in tech.*

↘ submit your questions to: editor@maximumpc.com

THE NEWS

The empire strikes back

Alder Lake is just the start for team blue

INTEL'S CHIP development might feel as if it has stalled a little compared with AMD's blistering development cycle, but the blue team hasn't been idle, and rumors about the forthcoming 12th generation Alder Lake chips have been coming thick and fast. After what seems like far too long, Intel will soon have something genuinely innovative.

Alder Lake will be a hybrid processor, combining Golden Cove cores for performance, and Gracemont cores for efficiency, with up to eight of each. It will use an LGA1700 socket with DDR5, and PCIe 5.0 support, but the accompanying chipset is unlikely to feature PCIe 5.0. What you plug into your PCIe 5.0 bus is another matter, there won't be any PCIe 5.0 SSD drives until later next year, and we've yet to see a consumer graphics card that can benefit from PCIe 4.0, let alone 5.0.

Rumor also has it that Intel will start at the top, with unlocked K, and KF versions, along with the performance Z690 chipset. We have seen leaked shots of a motherboard sporting 20 VRMs, so something wants a lot of carefully controlled power.

If this is correct, it means Intel is looking to make a splash from the start. Cannon Lake hardly made a ripple, and neither Ice Lake nor Comet Lake did much better. Even Rocket Lake proved somewhat pedestrian after the smoke had cleared. Meanwhile, each iteration of AMD's Zen landed with a bang.

A few early Alder Lake samples have been floating about, including a Core i9-12900K. Overall, the word is to expect a 10 to 15 percent performance bump per watt over Rocket Lake, with up to a 20 percent increase in single thread speeds. The first Alder Lake chips are due this fall.

Alder Lake will use the Intel 7 process node and, no, we didn't miss the 'nm' off, Intel has decided to stop using it. For decades now we've classified chips as using a process node, which uses the nanometer as its metric. You might think that chips built using a 14nm process will have 14nm features at least somewhere, but nope.

It used to be that the process number matched the transistor's gate size and the half pitch, but that hasn't been true since the late 90s. After that, it gets complicated. For a while the node matched the half pitch, but this too became detached. These days, there is no fixed measurement on a chip that equates to the given process node. The number is calculated differently by each manufacturer, a rough combination of measurements, and flimflam.

Different manufacturing processing, and lithography methods have muddied the waters further. The nm value is now just a marketing number: a new node simply means there's a new manufacturing process which is better, so needs a 'better' number.

Intel has invented its own process number, unattached to nanometers. Changes in the Intel node names will now be linked to major improvements in performance, power consumption, and density. When what Intel previously called 7nm arrives, it will be called Intel 4. In 2024, we hope to see Intel's 'breakthrough' 20A process using ribbonFETS.

There are two answers to this. One is that Intel is correct. The other, more cynically, is that it would say that, wouldn't it? Intel is, and has been for a while, losing the node marketing race. It spent years stuck at 14nm, as others raced to 10nm, and 7nm. Its 10nm chips haven't yet reached the desktop market, and while rivals are already looking at 5nm, it still sells 14nm chips.

This is slightly unfair though, Intel's 10nm is similar to Samsung and TSMC's 7nm process. You can understand why Intel wants to drop the nanometer as a metric. A similar move took the base clock frequency out of the chip name a few years ago.

Intel has also revealed some of it roadmap after Alder Lake, which will be followed by Rapter Lake, possibly as early as next year. It will also use the 600 series chipset, and looks to be a refresh of Alder Lake. Following that, there will be Meteor Lake, and a new LGA1800 socket. This is a tile design using Foveros technology, so Intel can mix, and match components. The company plans to be back on top with a commanding technological lead by 2025.

Of course, the blue team has big plans for the future. Alder Lake should put it back in the running, and it appears to finally have a solid roadmap behind it. Bring it on. **–CL**

> ❝
> **With the Alder Lake chips, Intel is looking to make a splash from the start.**

© INTEL

NEW ENERGY LAW HITS ALIENWARE
DELL ISSUES WARNING

NEW LAWS ON ENERGY EFFICIENCY have come into effect in California and five other States. The California Energy Commission Tier II implementation defines a mandatory energy consumption for desktop systems, depending on the PC's memory, GPU, extra drives, and so on. Warnings appeared on Dell's website under selected Alienware systems, saying they could no longer be shipped to six States. Closer inspection of the rules reveals they only apply to a PC when it idles, not under load, and the idea is to limit pointless power consumption, not performance. Despite having five years' notice, some configurations of Alienware's Aurora R10, and R12 systems managed to flunk the test, but as no other system builder had any problems, it seems our rigs are safe. **–CL**

WIN11 GOES BETA
Try it for yourself now

MICROSOFT HAS PUT Win11 build 22000.100 into beta on the Windows Insider Program, the first build to do so, and a sign that Microsoft is feeling confident that it's ready.

Previous builds were limited to the development channel. The Windows Insider Program may sound like a fancy club, but joining is easy—you just need a Microsoft account, and to tick a box to say you understand you might experience data loss, crashes, and security vulnerabilities (what's new?).

There are a few missing features, such as Chat on Microsoft Teams, and no doubt the odd bug will surface. The next step will be a version on the release channel of the Insider Program before it appears on new systems later this year. Microsoft also launched a free trial of its cloud-based Windows 365 system but suspended sign-ups just two days later. Despite only being available to businesses, it proved so popular that Microsoft was running out of server capacity. Or it was an astute publicity stunt? **–CL**

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

TRIUMPHS

✓ **3D PRINTED BRIDGE**
A stainless steel footbridge, weighing 4.9 tons, across a canal in Amsterdam has been made using 3D printing.

✓ **WESTERN DIGITAL BOOST**
Spinning discs are alive and well, with WD reporting sales were up by a third in the last quarter.

TEN BUCK PC
Some bright spark managed to get a $10 microcontroller (an Espressif Systems ESP32) to emulate an original 1981 IBM PC.

TRAGEDIES

✗ **HACKERS DUMP FIFA**
After failing to extort $28m from Electronic Arts, hackers who stole the FIFA 21 source code have dumped it online.

✗ **ORIGINAL KINDLES DYING**
As 3G networks are shut down older Amazon Kindles are left with no connection. Amazon is offering new ones instead.

✗ **REFORGED, NOT FINISHED**
Leaks show that Blizzard knew *Warcraft 3: Reforged* wasn't ready but released it anyway.

AMD'S RADEON 6600 XT LAUNCHED
AMD's releases schedule just keeps rolling

THERE'S ALWAYS SOMETHING cooking at AMD, and this month it's the launch of the Radeon RX 6600 XT. At $379, this mid-range RDNA2 card costs $100 less than the 6700 XT and is aimed at 1080p gaming. It looks expensive compared with the RTX 3060, which starts at $329, however, you won't find a 3060 at anywhere near that price at the moment.

Inside the 6600 XT is the new Navi 23 GPU with 2,048 stream processor, and 32 compute units running at a base clock of 1,968MHz, and a 2,359MHz boost. The GPU is significantly smaller than the Navi 22, which means more chips per wafer. That's good news for supplies, as small chips mean more chips. There's 8GB of GDDR6 and 32MB of L3 'Infinity' cache, quite a drop from the 6700's 96MB, but again, it helps with supply issues.

With the graphics card market in a strange place, AMD has concentrated on volume, and the result is an expensive 1080p card that's about 15 percent faster than the 3060. Unusually, there is no AMD reference design, but while the design may be a compromise, it should actually be available. We'll find out just how good it is next month.

AMD has just had a record-breaking second quarter, with year-on-year revenue jumping 99 percent. AMD's CEO, Lisa Su, claims it "remains on track to launch next-generation products in 2022". This means 5nm Zen 4 processors, RDNA 3 GPUs, and Ryzen 6000 series processors. We can also expect an assault on gaming and the high-end laptop market. With the first Ryzen only appearing in 2017, five generations in four years is a truly hectic pace. Long may it continue. **–CL**

PLASTIC CHIPS
ARM builds a flexible SoC

INTEGRATED CIRCUITS have been constructed on a silicon substrate since their inception (we even use the word silicon as a synonym for chips), but ARM has been working on using plastic as a substrate for about ten years. Simple electric circuits on flexible plastic are nothing new; paint a design in conductive ink, and you're done. There's one in every membrane keyboard and RFID sticker.

Researchers at ARM and PragmatIC have managed to put a fully-functional SoC on a flexible plastic substrate. It is essentially a Cortex M0 microcontroller. The PlasticARM (great name) carries 128bytes of RAM, 456bytes of ROM, and a Cortex M0 core that supports 32-bit ARM microarchitecture. The plastic component is a polyimide, and the transistors are thin metal oxide. It's made using a hybrid mix of a traditional photolithography process, and specialist printing. There are 13 material layers, with four metallic layers, totaling 18,000 logic gates in all.

The regular silicon Cortex M0 is built by TSMC on a 90nm process, and fractions of a square millimeter, the plastic version is just under 60. It's also slower than the original, running at kilohertz rather than megahertz. Energy efficiency isn't good either, as it loses a lot to static. Currently, it can only run programs hard-wired into it, but re-writable memory is coming. This isn't a rival to traditional silicon as such, but rather a chip for simple, mundane, or novel applications, where power and speed limitations aren't an issue.

There are no specific plans for the PlasticARM just yet, however, the goal is to make plastic chips cheap enough to be disposable. Eventually, they will be priced in cents rather than dollars and printed in their billions. Then you can have chips on anything, and everything. **–CL**

Game fries cards

AMAZON'S NEW MMORGP, *New World*, has been delayed until the end of September, due to a rash of technical problems. The game had an extensive beta test, with up to 180,000 simultaneous players, before reports started coming in of graphics cards crashing, and burning. For some unknown reason, EVGA RTX 3090 cards overheated to the point of destruction, while AMD cards appear to be immune.

Thermal management is usually good these days; drivers won't let software hammer cards to death, and there's hardware and firmware on the cards to stop things from getting too toasty, so it's unclear how much blame lies with the game. The launch of the RTX 3090 was marred by crashes, and the problem was traced to the type of capacitors fitted, but whatever the cause, it is clear that the game has highlighted a problem. Testers reported seeing a 100 percent load on their graphics cards during the loading screens though capping the frame rate reportedly cures it.

New World is the first major outing for the Open 3D Engine, an open-source game engine jointly developed by Amazon Games, and the Linux Foundation, so a lot is riding on it. However, a delayed game is generally a better game—it can be hard to rise above initial disappointment, as *Cyberpunk 2077* shows. The developers have said they will use the extra time to polish the "overall experience". **–CL**

The *New World* beta was literally too hot for some graphics cards to handle.

Facebook bans researchers into Facebook

Researchers looking into advert transparency, and the spread of misinformation on Facebook found themselves banned for a violation of its terms of service.

The group was examining the spread of political adverts on Facebook and had developed a browser plug-in to collect data on which ads are shown and to whom. This is a hot topic of course, as Facebook's political adverts aren't fact-checked. The company makes some data available, but not how the ads are targeted.

One researcher said, "Facebook is silencing us because our work often calls attention to problems". However, Facebook has had its fingers burned by third parties data scraping before. **–CL**

Google builds a time crystal

There is, so the theory goes, a new state of matter, along with gas, liquid, solid, plasma, and so on. A time crystal is an object caught in a repetitive motion that does not lose energy, which not only sounds impossible but also breaks the second law of thermodynamics.

Time crystal was first proposed in 2012, but now Google scientists claim to have used a quantum computer to demonstrate that it is possible. It's also a convenient advert for Google's quantum computer, called Sycamore, which is capable of making calculations that would take conventional computers millions of years.

If this quantum computer has indeed helped prove the possibility of the time crystal, it has earned its development budget. **–CL**

Jarred Walton

TECH TALK

Prepare for the DDR5 shift

BUCKLE YOUR WALLETS, because the Intel 12th Gen Alder Lake and AMD AM5 platforms will entail some costly upgrades. Modern PCs utilize DDR4 memory, but that's all set to change with the next generation CPUs and platforms, ushering in a new era of memory bandwidth and capacity. The last time desktop PCs went through such a transition was five years ago, and early adopters often pay the price. Here's what to expect from DDR5.

DDR5 slots aren't compatible with DDR4 slots. We might see motherboards that support both DDR4 and DDR5, but you'll have to choose between one or the other—you won't be able to install both memory types at the same time. For this reason, I've avoided mixed DIMM slot motherboards in the past; I'd rather have four of one type of slot rather than two of each. DDR5 slots will maintain the same pin count as DDR4, 288-pin, but the pin layout will be different, and the slots are keyed with a notch so that you can't put in the wrong type of memory. Voltages are also different, with DDR5 dropping to a standard 1.1V compared to 1.2V on DDR4.

DDR5 reorganizes the memory chip layout too. DDR5 supports up to 32 memory banks distributed over eight bank groups, compared to 16 banks and four groups on DDR4. The burst length for DDR5 is also doubled to 16, from DDR4's eight, and DDR5 includes a Same Bank Refresh Function (SBRF) that allows the memory modules to use a different bank to boost throughput.

Besides improved efficiency and throughput, DDR5 boosts the clock speeds. The official JEDEC standard for DDR5 starts at 3200 MT/s, matching the fastest official DDR4 JEDEC speed. Maximum official speeds extend to 6400 MT/s. That means a typical dual-channel desktop or laptop memory configuration will have a maximum bandwidth of 102.4 GB/s—something that should

help improve integrated graphics performance. There will inevitably be overclocked DDR5 chips that reach higher speeds. The fastest DDR4 kits can break 5000 MT/s already, and there are already plans for at least DDR5-8400.

It's not just speed, though. Capacities are set to get a massive boost. DDR4 chips right now top out at 16Gb, which means the typical DIMM with eight chips on each side has a maximum capacity of 32GB. DDR5 is set to have chip capacities of up to 64Gb, meaning DIMMs could have up to four times the capacity—128GB DIMMs, coming in the next year or so! DDR5 also supports die stacking with up to 16 dies on a single chip for servers, which means a maxed-out Load-Reduced DIMM (LRDIMM) could theoretically have 4TB of capacity. Forget about cost—just imagine that. A server with 16 DIMM slots could potentially run 64TB of RAM.

There's more to life than just speed and capacity, however, and this is where we have to hit pause on all the excitement. There are many applications—including games—that respond better to reduced memory latency rather than more bandwidth. Official JEDEC standards for

DDR4 have timings of 20-20-20 for DDR4-3200 as the best speed, though enthusiast kits with 14-14-14 timings exist. That's 12.5ns of CAS Latency for CL20 and 8.75ns for CL14 DDR4-3200.

The fastest official DDR5-3200 timings now are 22-22-22, which means 13.75ns of latency, and the fastest DDR5-6400 memory has 46-46-46 timings, or 14.375ns latency. We'll see what real-world timings end up shipping, but DDR5-3200 kits that adhere to JEDEC standards could perform worse than current DDR4-3200 kits in situations where responsiveness and latency are more important than raw bandwidth.

The same thing happened with the transition between DDR3 and DDR4 memory, so there's nothing new about these concerns. Over time, we expect even 'normal' DDR5 kits will end up outperforming enthusiast DDR4 kits, just like most DDR4 kits now use unofficial timings—a quick look online shows that even the cheapest DDR4-3200 kits are CL16. But it will take time to reach that point, and there will undoubtedly be a transition period where DDR5 kits cost significantly more than equivalent performance and capacity DDR4 kits.

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

> SK Hynix will offer a range of DDR5 solutions, such as the 64GB DDR5-4800 DIMMs.

THE LIST

MAXIMUM PC'S 8 UPCOMING GAMES OF 2021

8

DEATHLOOP Break free from a time loop while being hunted down in Arkane Studios' action-shooter.

7

S.T.A.L.K.E.R. 2 - HEART OF CHERNOBYL Intense survival sequel will keep you on your toes.

6

SABLE Glide across stunning deserts and plains, exploring ancient ruins in a beautiful open world.

5

TOTAL WAR: WARHAMMER III Four years on, a third (and final?) iteration to look forward to!

4

HALO INFINITE Infinite's fast-paced FPS gameplay will have you hooked (yes, that's a pun; there's a grappling hook now).

3

FAR CRY 6 The franchise now takes us to Yara to engage in some modern high-octane guerrilla warfare.

2

FORZA HORIZON 5 Fun, without being OTT, racing in an open-world Mexican landscape? We're in!

1

BATTLEFIELD 2042 An epic shooter set in a crisis environment caused by natural disasters.

© BETHESDA SOFTWORKS, GSC GAME WORLD, RAW FURY, FERAL INTERACTIVE/SEGA, XBOX GAME STUDIOS, UBISOFT, MICROSOFT STUDIOS, ELECTRONIC ARTS

Zak Storey

Zak is Maximum PC's
editor-in-chief and
long-time staff member.

TRADE CHAT

Microsoft Could Threaten our National Security

WINDOWS 11 IS PERHAPS one of the most divisive operating systems we've seen come out of Microsoft for some time. It's not any one design element that's causing the problem, the biggest bugbear, by far, is the security requirements—and, oh boy, does it open up one heck of a can of worms.

So, to run Windows 11, Microsoft has declared that you're going to need a processor that supports TPM (Trusted Platform Module) 2.0, along with Secure Boot. Now, I won't go into detail about what TPM is, or how it works, as we've covered that enough already. However, the short version is that this is usually used in more enterprise grade products and companies than with everyday users, and it deals with generating and storing cryptographic keys, it's a crypto-processor effectively.

The problem, however, is that Microsoft is demanding TPM 2.0 at a minimum to run Windows 11, despite it only being available on Intel 8th series processors and AMD Ryzen 2000 series chips and above. If you haven't built a machine in the last three years or so, you won't be able to upgrade.

On the one hand, I can understand why Microsoft is pushing for this. It's trying to take the fight to Apple in regards to the security of its operating system (in fact there's a myriad of new security features baked into the update, that make it plain to see). However, there are some massive drawbacks to this too, that absolutely need addressing and resolving before it goes live.

First up, we know that Windows 10 has an EOL date now, and that's October 14th, 2025. That's just four years away, admittedly a lengthy amount of time, but it will mark the end of all security updates for the supposedly perennial platform without any intervention from Microsoft.

Enhanced features can't hide the major problems Windows 11's security demands will create.

Why does that matter? Well, my biggest concern isn't actually to do with us as individuals at all. Inevitably in those next four years, the majority of the systems we're running today, as enthusiasts, will need to be upgraded anyway. Those who do want to try Microsoft's latest and greatest will probably already be looking at updates, or discrete TPM 2.0 modules to install on motherboards instead, so it's less of a problem there as well.

However, the bigger issue is for health services, government departments, and infrastructure companies around the world. The majority of which typically do favor running cheaper, older, machines.

If you're asking all the big players in Western society to upgrade the entirety of their digital infrastructure to more secure devices, built within the last three years, so they can get security updates, that's an enormous cost. You're going to need hundreds of thousands of new machines, bespoke software will have to be rewritten, and bugs quashed. You will also need to retrain your staff and, of course, we're in the middle of a processor shortage, too.

The environmental impact will be huge, and the financial implications massive, which will inevitably be pushed on to taxpayers and customers. But if you don't do it, the basic infrastructure of your society could be susceptible to ransomware attacks and worse, which we've seen only proliferate in the last decade. All of which poses a threat to national security.

It's a genuine nightmare waiting to happen, and perhaps the most asinine decision the company has ever made. Don't get me wrong, I love using Windows 11, it's a fantastic operating system. The new features are amazing, and I can't rate it highly enough. But this decision to demand TPM 2.0 at an absolute minimum, makes the world and all the countries that use Windows as a platform, less secure, not more.

> **If you don't upgrade, the basic infrastructure of your society could be susceptible to attack.**

DOCTOR

> ## M.2 cooling
> ## Strange boot issue
> ## 2FA backup

Cooling M.2 drives

I've spent several weeks seeking advice on this topic, but haven't found a spec for the recommended operating range for M.2 drives. All I've found are mixed messages suggesting they do need cooling and they don't need cooling. Folks have asked the question because they were concerned after seeing high temperatures via monitoring software like Speedfan and CrystalDiskInfo.

There's also a question over whether to purchase cooling with a fan or without. My first attempt was to mistakenly purchase an EZDIY=FAB 5V ARGB M.2 Heatsink for my Gigabyte Z370 HD3 motherboard, which doesn't have a 5V location to plug it in. Instead, I've ordered a Sabrent M.2 2280 Heatsink that has no fan. One final question: do you recommend using thermal compound paste instead of the "thermal pads/tape" they come with? **–Dean Johnson**

THE DOCTOR RESPONDS:
Quite a bit to unpack here. The simple answer is: unless your drive is running consistently 'hot', there's no point in fitting any cooling. But what constitutes 'hot'? First, M.2

Monitor your M.2 drive temps before plumping for a heatsink.

drives naturally run hotter—by around 20°C if the Doc's PC is anything to go by—than other installed drives. That's down to where it's installed on the motherboard, typically near the GPU and CPU where temperatures are higher.

M.2 drives supply two temperature sensors—one for the drive itself, and the other for the controller. The latter is consistently hotter than the former, but it's the former you should monitor closely. CrystalDiskInfo and your drive's official tool (such as Samsung Magician) will provide a real-time display of the drive's temperature, but you can monitor both sensors using a tool such as HWiNFO64 (www.hwinfo. com). It provides minimum and maximum figures to go with the current setting.

The M.2 drive's primary sensor on the Doc's PC reports the drive holding a steady temperature of around 50-55°C, while the controller temperature is around 15°C hotter. This is well within acceptable parameters. You only need to start worrying when the drive reaches 75°C or higher, which could induce thermal throttling, reducing the drive's performance to ensure it stays within temperature limits.

In the vast majority of cases, your drive isn't going to be taxed enough to drive its temperature to unsafe levels except in limited circumstances, and that's when the throttling should kick in.

The only time you'd want to consider a heatsink is if those idling temperatures are much higher than the comfortable norms observed on the Doc's rig. That might indicate airflow issues inside a cramped case, and it's likely your other components will be running hot too. Start by addressing any obvious airflow issues such as poorly stowed cables or clogged-up fans. If your mobo has two M.2 ports of identical spec, you might want to consider moving the drive to the other one to see if runs a little cooler there.

If you're unable to bring down the temperature of the drive without external help, then invest in a heatsink such as Darren's Sabrent M.2 2280 ($25-30, www.sabrent.com). Note that if you do that, you should look to remove any labeling covering the NANDs to improve heat transfer to the heatsink.

As for applying thermal paste, the overall effect on reducing temperature will be minimal, so the Doc would recommend sticking with the thermal pads that ship with the heatsink.

↘ submit your questions to: doctor@maximumpc.com

Windows boot errors

Hey Doc! I just built a PC with ASUS ROG Strix B550-F gaming mobo, Ryzen 5 5600X CPU, and 1TB M.2 NVMe drive for Windows and apps. It replaced my old overclocked i7-4770K machine and I'm pleased with the results with one major exception: whenever I make an adjustment in the UEFI, Windows won't load, and I receive a message that the \system32\winload.sys file is missing. I then click 'Select an OS', choose the only option (Windows 10), and then after an hour of flashing disk lights and dark screen Windows finally loads. I run ChkDsk and a few errors are noted as having been repaired. The system then runs okay, but I'm avoiding any BIOS adjustments. If I then go back into the UEFI and verify the M.2 drive is the first boot device, and then exit without saving, the PC boots without problems. Can you explain why this is occurring and if there is anything I can do to make BIOS adjustments reliable?   **–Philip Wiener**

THE DOCTOR RESPONDS: There were lots of questions we wanted answers to before we could suggest anything concrete, but by the time Philip replied to our email, he'd already found – and fixed – the source of his problem. Prior to now, he'd routinely moved his paging file off the system drive – in the case of his new PC, from the M.2 system drive to his 2TB SSD. Restoring the paging file to its original location fixed his problem, and he can now boot into Windows in around 20-30 seconds.

The old trick of moving your paging file to a separate drive to improve performance is pretty much redundant these days with fast M.2 and SSD drives, particularly as Windows is likely to be installed on the fastest drive in your system. To remove any manually set paging file settings, right-click the Start button and choose System, then click 'Advanced system settings' to open the old System Properties Control Panel at the Advanced tab. Click Settings... under Performance, switch to the Advanced tab, and click Change... under 'Virtual memory'. Tick 'Automatically manage paging file size for all drives' and click OK three times, rebooting at the end of the process.

App Store follow-up

I just read the question from Gerry Exstein about not being able to download the calculator from the Microsoft Store. I have, through my own experience discovered why when you hit 'Get' nothing happens. Unknown to me, Microsoft limits the number of devices a logged-in user can use to access the Store. Officially it's ten, but in many cases – including my experience – it tops out at five. I had this issue when I got my Surface. I struggled for hours then discovered this limit. I went into my MS account to devices and saw only six were registered, including a few I no longer owned. I deleted these redundant devices and was able to use the Store without any problems!
                    –Gregory Watts

THE DOCTOR RESPONDS: They say you can't teach an old Doc new tricks, but Gregory's helpful tip proves the opposite, so thanks for sharing it. To see what devices are registered to your account, open your browser and sign in at https://account.microsoft.com/account. Click the Devices section to view a list of devices currently connected to your account. Click a redundant entry and you'll see a 'Remove device' option. Click this to free up slots in your Microsoft account and ensure you don't fall foul of the mysteriously malfunctioning Store problem.

Protect 2FA codes

Hi Doc, I'm looking into switching on two-step verification for as many

Authy can back up your 2FA codes securely to the cloud.

online accounts as possible but am worried about what happens if I lose access to the 2FA codes (say, if my phone was stolen). What options are there for me to securely back up these codes?   **–Steve Cox**

THE DOCTOR RESPONDS: When you come to set up a 2FA code with an online account, you should be given an option to generate a 16-character 'recovery key', which you can use to recover the 2FA code on your new device through a new instance of your 2FA app. You can also generate recovery keys for other purposes too, for example, to recover a Firefox Account when you forget your account password.

You can print this recovery code out or write it down, or copy and paste it somewhere secure, say inside a Bitwarden secure note. Of course, you're trusting Bitwarden's servers to securely host your passwords in this instance. If you don't want to do that, check out the December 2019 issue for our guide to running a self-hosted Bitwarden server.

What can you do if you haven't generated these codes? First, log back into the account in question and navigate to the security section to see if you can generate a new recovery code for the account or 2FA code; if you can then problem solved, but if not, you can disable 2FA for that account, then set up a new 2FA code, making sure that this time you do take a copy of your recovery key.

If this seems like too much effort, then another option is to switch your 2FA app to Authy (https://authy.com). Authy allows you to sync your recovery codes between as many trusted devices as you like—phone, tablet, computer (via desktop app), and browsers (via add-ons). And if that's not sufficient, Authy also supports secure backups to the cloud via its servers—generate a strong password and your codes will be online and retrievable should you lose your device. Both sync and backup options are optional and can easily be disabled depending on what method(s) you trust.

TPM restart issue

I read that you were able to help a guy fix his strange shut-down issue and I thought about the intermittent problem I've had with my own system: built around an MSI Z170A GAMING M5 1.0 (BIOS v 1.F0). From the beginning, the machine would go directly into the BIOS whenever I attempted to restart (as opposed to shutting down and starting up manually). I exit without saving and the machine boots fine 100 percent of the time. Occasionally for Windows Update, the machine will restart properly, but not for any other program or if I try to restart from the Start menu. I've found nothing in my searches. Can you help?
                    –Zach Pulkin

THE DOCTOR RESPONDS: Before we had a chance to speak to Zach, he stumbled on the solution to his problem while reading about Windows 11 system requirements. By enabling TPM 2.0 on his board, his PC now restarts properly. He asked us to print the tip for others, which we're obviously happy to do. ⏻

Rise and shine, pull back the curtains, there's a new Windows in town.

GENERALLY SPEAKING, there are two things that an update can bring to the table. One of which is an element of fear and the other, an element of excitement. We fear that features we have grown to love and use every day may vanish (we haven't forgotten about you Clippy). We also fear that too much change will interrupt the ease of use we have built up by using a certain OS for years on end.

This fear of the unknown can make you hesitant about updating. Certain parts of an operating system make you fall in love with it, especially if it helps day-to-day life, whether that is work or just domestic social use. Also, new bugs can be introduced and it may take a while for everything to get back to running perfectly, so it may seem like a step backward initially after updating.

On the flip side of that, the unknown can bring excitement, especially when you have been faced with the same layout and overall design consistently every single day for many years. It's nice to see an overhaul where new design elements are brought in to spruce it all up. Other features that help efficiency and overall bug fixes can make your PCs run smoother and better too. These are always welcomed with open arms and a new update can often sort out an underlying issue on your PC.

Microsoft's latest OS, Windows 11, is being released in October this year. Hopefully, it should be a simple upgrade for many people (as long as it isn't a tablet OS like Windows 8). However, with some major compatibility issues and an end-of-life date set for Windows 10, it has already kicked up a storm. So let's dive straight into this new OS and get a feel for what it's made of. –SAM LEWIS

WINDOWS 11 PRIMER

FIRST VISUAL IMPRESSIONS

1. ANDROID APPS?

Yes, you did read that correctly, and no, we haven't lost the plot just yet. You can get android apps on Windows and not through some dodgy workaround. The Microsoft app store has had a little bit of tweaking, and though it isn't too different visually, it now features rounded corners, which makes everything a bit easier on the eye.

One of the main positives is that, regardless of the streaming service, it can now show movies, TV shows, and more. This is a pretty neat feature in itself, however the mic drop news is that Amazon's Android app store will be integrated into the Microsoft app store. That's right, you can now get a designated TikTok app running on your PC, if anyone wants that.

Sarcasm aside, this is great news as the ghost-town-like app store of

STRAIGHT OFF the bat, you can tell that this new OS has had a fairly hefty re-design, and hats off to Microsoft for this one. The new design is incredible, it takes away the sharper lines from Windows 10 and softens them down. It has a frosted glass effect with plenty of 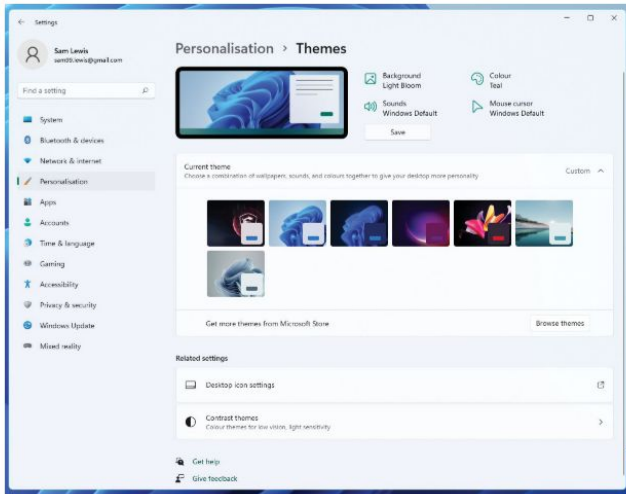different themes and color palettes to choose from. This change creates an easy-on-the-eye UI. Although there were plenty of customization options in Windows 10, it still felt like it had a harsher contrast and felt too utilitarian at times. Windows 11 brings a lot more color to the interface without looking too playful. Even though most people will change these out regardless, the new wallpapers add a bit more pop to everyone's first glance at the new OS, which is a nice touch.

There are certainly some major similarities with another huge competitor, lets just say, we are getting a lot of MacOS vibes from the new look. The biggest elephant in the room has to be the new centered dock. As a whole, this has gotten bad press due to the undeniable link to its Apple rival, but it doesn't just stop there. Windows 11 introduces a simple change that affects all areas of the new UI and that is rounded corners. Yes, Apple doesn't own the rights to rounded corners but they are certainly synonymous with softer shapes and a more elegant design scheme, which can now be found in Windows 11. The similarities here are uncanny, but it's not an OS clone and still operates in the same format as the previous OS. One more noticeable thing that Microsoft has as a counterpart to a macOS feature is that Microsoft Teams is now built-in, kind of a Facetime iMessage alternative, right? If you don't use it, you can always uninstall it, but it's nice to see Microsoft integrating a great communication tool into its new OS. Anyway, enough with the Apple comparisons for now.

All design is subjective, so if you prefer the older design in Windows 10 you won't be alone, as there are plenty who equally like and dislike this new step forward for Microsoft. If you are a Mac user and have always wanted to get into Windows there hasn't been a better time than now to transition over. Due to the recent overhauls, it will feel a lot more like home to those who do that.

However, the centered dock alteration has led to a lot of people running around screaming with disgust considering this is a big change for Windows. Luckily, and without installing another piece of customizable software, you can revert the centered dock to the left. Simply head to settings, personalization, taskbar, and then go into the taskbar behaviors tab. Here you will find taskbar alignment, where you can set the alignment to the left and end those sleepless nights. Don't be put off by everything you see here, if anything this latest version of Windows is the most customizable yet and an exciting step, design-wise, for the software.

5. SNAP LAYOUTS

If you hover over the maximize button in any window, a pop-up will appear with a few options for laying out your open tabs. It's an efficient way to organize your work or make the most out of that ultra-wide monitor. If you use your PC for creative work, it makes perfect sense to take advantage of

© MICROSOFT

yesteryear will now be packed with plenty more applications. This is generally going to benefit tablet or touchscreen devices, as this is what most Android apps are designed for, but nevertheless, it's nice to have the option there.
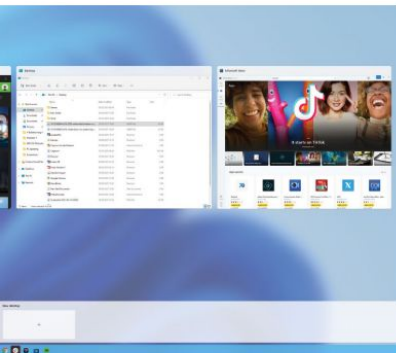
this new feature. Although snapping windows together isn't necessarily new, it's a much easier method. You could technically create these layouts before with other software or by just dragging windows out to fit, but this gives you templates to choose from that make full use of the screen. It's a great built-in feature of Windows 11, but you can disable it if you want.

2. WIDGET PANEL

Eagle-eyed readers may have spotted an icon towards the bottom of the screen that looks like a window, this is the new Widgets panel. Think of it as an updated version of Windows 10's news and interests panel. It shows widgets at the top followed by personalized news and information below. This sits in a neat panel-like format with a frosted glass background and is similar to when you are using a tablet or touchscreen Windows device, as you can swipe from the side to open this up.

Widgets include weather, sports, to-do list, calendar, photos, and traffic updates. Nothing breathtakingly new here, it's just all laid out efficiently and is easy to access at the click of a button. There's even a web search bar at the top so you can get procrastinating.

3. TASK VIEW

There is another newly designed icon in the taskbar, the task view. Compared with Windows 10, the newer OS seems to have taken a less-is-more approach. At first glance, everything is more evenly spaced and easier to follow, more intuitive, and smoother—especially the new animation that opens up the task view. It spaces out open tabs evenly, without cramming them all together so you can quickly transition between them. When the task view is open, it frosts the background, which seems to be a theme in the new OS. At the bottom of the screen, you can see your open desktops and easily open up a new one. You might already be sick of us saying this, but everything is rounded and it looks neat and tidy.
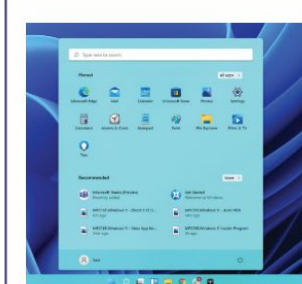
4. START MENU

Last but not least the start menu, the main hub for Windows 11, has had a big visual change but is still easy to navigate. Live tiles have gone and now we have a smaller pop-up window with, you guessed it, rounded corners, and frosted background. At the top of the window is pinned apps, which is a great place to keep all your most important applications. Next to this on the top right, there is a button for all apps and when you open this it is a familiar list format like in Windows 10.

In the bottom half of the start menu is a new recommended icon list that can display recent documents, applications, and files you have opened or frequently visit. This is handy if you close down things when you don't mean to. The reasoning behind Microsoft's new start menu is that the user can get to what they need to quickly. That's really one of the main findings from the new OS, it's less complex, cleaner, and more intuitive. Sounds good right?

ISSUES ARISING
TO BE (COMPATIBLE) OR NOT TO BE (COMPATIBLE)

IT ALL SEEMED TOO GOOD to be true right? Well in fairness it is very good, and we should all be excited about its full release. However, the elephant in the room is a big one at that. You may want to get hold of Windows 11 but can you actually do that? The likelihood is that as of now, it's not compatible at all. If your PC is over 3 years old, the odds are that it might not be compatible with Windows 11 and that is a worrying thought. Don't however be too comfortable continuing with Windows 10 if you can't upgrade as that now has an end-of-life date in October 2025. Even if you have a fairly new PC don't be too relaxed as it still may not be compatible. Yikes, and yikes again for good measure.

Windows 11 is the first new version of Windows since 2015 and a major part of this new software is to do with its upgrade in security. For as long as we can remember, security online has always been an issue. So much valuable data is stored and accessed on our computers daily and this is something that we need to keep private and secure. There are plenty of different ways to protect ourselves but there is no promise that our data will not get leaked or our systems are totally secure. Of course, VPNs help and the same goes for anti-virus software too but it's always nice to know that your OS is secure seeing as this is the foundation of your whole PC.

On Microsoft's website, it states that "Windows 11 raises the bar for security by requiring hardware that can enable protections like Windows Hello, Device Encryption, virtualization-based security (VBS), hypervisor-protected code integrity (HVCI) and Secure Boot". The company also states that "the combination of these features has been shown to reduce malware by 60 percent on tested devices". Now on paper, this sounds great, who doesn't want ramped-up security, the peace of mind alone is rewarding enough. However, one of the main points to take away from this is that it requires specific hardware to enable these protections. This hardware comes down to the CPU and motherboard. The CPUs supported by Windows 11 must have an embedded TPM 2.0 (Trusted Platform Module), must support secure boot, and support VBS and specific VBS capabilities.

Windows has had a long history of being attacked by viruses, malware, and security issues so we understand why they are trying to pump up the protection for its users, but by doing so it forces the consumer to update their perfectly good systems just to have the latest software. It seems like a way to flush out older less secure hardware. To run Windows 11, according to Microsoft, there is a pretty short list of compatible hardware. As far as minimum requirements go it needs to be a chip greater than a 1GHz dual-core 64bit processor, 4GB memory, and 64GB of storage. Not too intolerable right? It's when you start diving into the specific CPUs that are compatible the list becomes narrower. To run Windows 11 you will need an Intel 8th gen processor, AMD Zen 2, or a Qualcomm 7 or 8 Series or higher. Yeah, not exactly great news to a lot of people. To be fair, as they are releasing the pre-release software to Windows Insiders, they will be testing to identify if devices running on Intel 7th gen chips and AMD Zen 1 chips meet their principles and are compatible. Fingers crossed.

SECURITY THREAT
Again though, you need a TPM 2.0 chip on board. This is a security chip that is either built into your CPU or motherboard and used for things such as encryption, increased security, and password protection, for example. Most systems of the last few years will have this as standard but do your research and check whether your system has a TPM 2.0 chip.

These compatibility issues could cause a lot of problems and a big backlash for Microsoft. Not everyone can afford to just switch up their rigs, it may even create a high demand for hardware embedded with TPM 2.0 and jump up the prices. This could have a negative impact on the resale market, just like how cryptocurrency has boosted the resell prices for GPUs to unfathomable levels in the past year. The last thing anyone wants is for a spike in CPU prices too, what next, PC cases reselling for over 1000 bucks? No thank you. If this actually happens please don't blame us, we'd appreciate that.

© ANDREW BROOKES/GETTY IMAGES

WINDOWS INSIDER PROGRAM

Microsoft is in a bit of a predicament here. The damage could already be done without a full release. This news of a limited amount of hardware being compatible with the newest update could have already pushed a lot of consumers away. Nobody wants to be forced to change what they have and be made to purchase something new in order to get the latest OS. On the flip side, hardware that is too old is unreliable and might not be powerful enough to take full advantage of the new software, and Microsoft will want everyone to have an equally smooth experience. It just seems because of the required TPM 2.0 chip there will be perfectly powerful enough PCs that can run games at 4K yet won't be compatible with Windows 11, which seems a bit unfair.

The compatibility problem isn't ideal, but hopefully, people won't bypass this update like Windows 8. We hope Microsoft takes all of this into consideration and can widen the compatibility of their newest operating system.

One of the best ways to get a taster of the latest preview builds of Windows 11 released by Microsoft, the Windows Insider Program is a massive community of people who sign up, run previews of the platform, then give their feedback to Windows engineers to help the early builds run smoother. With these builds, it's a good way to get them out to enthusiasts who can dissect the latest OS and discover any potential problems or bugs.

Signing up for the Windows Insider Program before the general release of Windows 11 is a great way to test it out for yourself and see whether you like it or not. It's also not that difficult to become a member of the Insider Program. Head over to the Windows Insider Program page on the Microsoft website and register by using a Microsoft account or creating one. The registration is free and super simple.

The next step is flighting (running the Insider Program builds). Once you are joined up to the program you have to choose a channel and then you can install an early version of Windows based on the channel you select. This channel will determine which updates you get. There are three to choose from: Dev, Beta, and Release Preview Channel. Think of these channels as difficulty levels in *Call of Duty*: Veteran, Hardened,
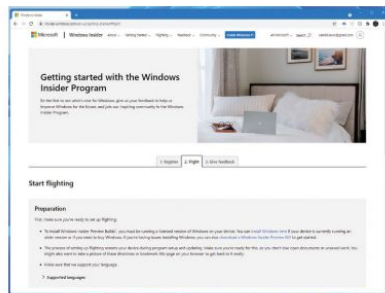
and Recruit, and it kind of makes more sense.

The Dev Channel (veteran) is, believe it or not, aimed towards developers or highly technical users and enthusiasts. This way the people who choose this channel will see the latest updates that come out. Nice, right? However these aren't always the most stable of builds so enter at your own risk. You might discover key problems and have to work around issues until they are resolved. If you provide feedback it can help the engineers and change the latest code.

The Beta Channel (hardened) is better suited for early entrants to the Insider Program who want to see what is next and upcoming for Windows. The difference here is that this channel is for people who want more reliability so they won't get the absolute latest but they will get reliable updates validated by Microsoft to ensure a more secure build of the OS.

Finally the Release Preview (recruit) the same applies, you'll get the upcoming version

of Windows before it's officially released. With this, it is made even more secure and stable with certain key features and advanced updates.

Once you have decided which channel is most appropriate, head to the Windows 11 flighting installation section on Microsoft's website and follow the instructions to update and install your build. If you are having compatibility issues with TPM 2.0, but you have TPM embedded already you might have to head to your BIOS.

Of course, all BIOSes are different, but the approach should be similar. When your PC is booting, mash the delete key to open the BIOS. Once in the BIOS, you need to find the TPM settings, hopefully, you can find this in the search function. For our build, we had to change the TPM device selection from Discrete TPM to Firmware TPM. Once you have changed your TPM settings, you should be able to restart your PC and follow the flighting instructions to get the latest build of Windows 11 to test and play with.

WINDOWS 11 GAMING FEATURES

PCs AND GAMING go hand in hand—unless you are using a notebook and trying to run *Cyberpunk 2077* at full settings at 4K, maybe. Anyway, millions upon millions of people game on their PCs daily and the gaming scene built up around PCs is fundamentally massive. Yes, console gaming has been around for ages too, but it has always suffered from hardware and software limitations. With PC gaming, the world is your oyster—customization is nearly endless, you can build a rig exactly how you like it and get it to perform and behave just the way you want (as long as you can get the parts you need).

Microsoft has been bridging its gap between Xbox and Windows closer and closer over the past few years for better integration. The tech giants are now aiming to support PC gaming further improving performance and accessibility.

1. DIRECT STORAGE

A feature that was first introduced on Xbox Series X and S, called direct storage, is coming to Windows 11. This API greatly improves asset generation in-game, enabling much larger draw distances and should help stability in gaming scenarios. It does all this while cutting down on CPU utilization and load times. Again this sounds pretty good, right? There are a few requirements to take advantage of this storage technique, this is more for future-proofing than anything. You will need an NVMe SSD with a standard NVMe express controller driver, shader model 6.0 support, and a DirectX 12 compatible CPU. If you have all this, the games will need to be designed to support direct storage, and while not every game will be able to take advantage of this, the feature should ensure your gaming is future-proofed.

2. AUTO HDR

This is another pretty neat feature that is seen making its way over to Windows 11 from Xbox Series X and S. Essentially, auto HDR adds high dynamic range lighting to games even if they don't support HDR. Yes, that gets a round of applause from us!

This feature benefits older titles, giving them a new light and making them look a whole lot fresher. For nostalgic reasons only, this is nothing but a positive. It gives these older titles a wider color range and brightness than before. Therefore the visuals are being improved without needing the help of developers.

The new OS will support auto HDR for titles that are developed for Direct X11 onwards. That means more than 1000 titles will be supported by auto HDR, and though there isn't a full list of compatible games yet, the likes of *Skyrim* and *Doom* will be supported. However, to take advantage of this feature, your monitor must already support HDR, of course.

3. DIRECTX 12 ULTIMATE

This is the latest graphics API (Application Programming Interface) by Microsoft and will be included in Windows 11. Included within this interface, there is a decent list of features. DirectX 12 Ultimate boasts Mesh Shaders, Sample Feedback, DirectX Raytracing, and Variable Rate Shading. These features are put in place to future-proof next-generation games. Ray tracing for current-gen titles is currently one of the biggest hype-words in gaming and rightfully so, it's pretty impressive.

DirectX ray tracing takes this a step further. Depending on current in-game demand, it will make ray tracing shaders load more efficiently. Scenes with fewer shaders will rely on inline raytracing to
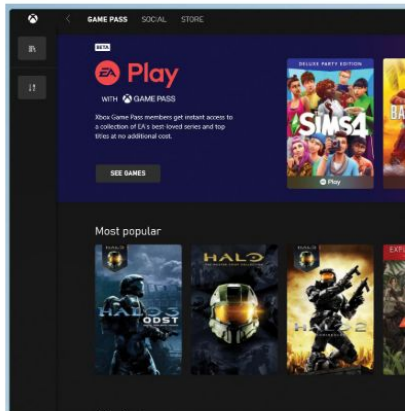
4. XBOX APPLICATION INTEGRATION

In Windows 11, the Xbox app is being integrated directly into the new software. This has many benefits, including creating a more seamless experience for gamers in the Xbox ecosphere. Accessing Xbox game pass for PC will be a much smoother process. The extensive games library already available on Xbox games pass makes this an exciting new addition to Windows and, with more coming to the service at launch, this is a feature we can appreciate. With Windows also easing up on peripheral compatibility restraints, you will be able to exploit this to play console games exactly how you like with your custom PC builds.

© IGN

THROUGH THE WINDOW
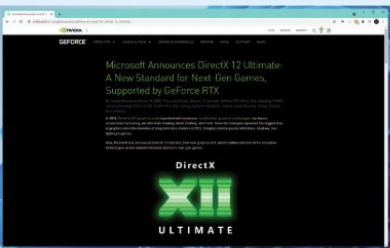
improve performance too. Variable-rate shading optimizes other scenes for improved performance, while it also leverages the GPU depending on the requirement necessary, therefore working in a variable motion.
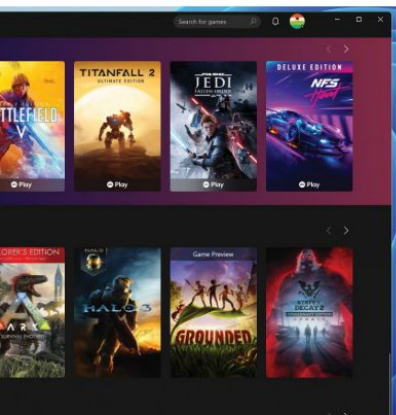
WINDOWS IS FOR GAMERS
Microsoft has indicated that it is here to support gamers and that Windows 11 is a place where gamers can fully enjoy their beloved titles by getting the most out of their games. It brings together a lot of elements that vastly improve the overall experience when it comes to games. By automatically integrating these features into the OS, it shows that the company cares about the impact gaming has on Windows. Where the compatibility issue sneaks in again is that, even if your PC is compatible with all of these new features, if it doesn't have the TPM 2.0 module, it still won't work with Windows 11. All of these features then become irrelevant, which isn't the best situation to be in.
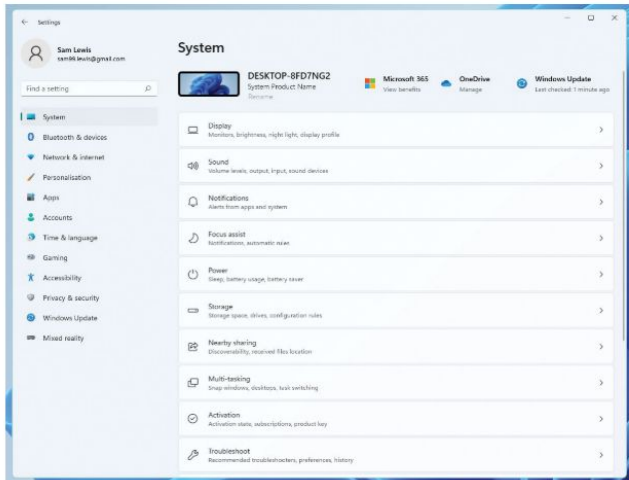
SO HERE WE ARE, Windows 11 is looming around the corner and we have been lucky enough to play around with Windows Insider builds to get a taste of what is to come. After the backlash it has already received over the compatibility issues, Microsoft has a lot riding on this full release. Will they distinguish these issues and widen support for older hardware? Will they be able to turn attention to all the great features that are coming out? We don't yet know, but with more people getting into the Windows Insider Program, hopefully, devices running on Intel 7th gen chips and AMD Zen 1 chips will meet their principles and are compatible. Yet again, a tighter security system protects users and Microsoft itself from cyber attacks, phishing, and data leaks, so every cloud has a silver lining.

We are impressed with the new look and feel of Windows 11. It's a beautiful new style and that new lick of paint has really freshened up Windows and made it a whole lot more exciting. With a design that is softer, brighter, and with a frosted glass aesthetic, the whole OS looks slicker, which is exactly what Windows needed. Paired with smooth animations, the whole feel of the operating system is extremely fluid. Subconsciously, these add up and create an environment for the users that makes everything easier, increasing productivity and workflow. We like the different themes too, everything seems to be over-engineered and for a design, this is what you want to see. The settings menu is another great change in this OS with a nicer format than in previous operating systems and also a lot more customization available. It displays everything a lot clearer than in Windows 10, a problem in the previous OS was that you could get lost in the settings, everything looked empty. Luckily with the new design language, this isn't an issue.

The anticipation for Windows 11 to succeed is nerve-wracking, due to the Windows OS curse of good, bad, good, and then bad. Windows XP will always be a nostalgic gem, but the disappointment of Vista that followed didn't sit well with a lot of people. After Vista, we were blessed with Windows 7, which remained strong until Windows 8 came along and focused more on touchscreen and tablet usage. Windows 10 brought clarity to Windows users once more with a solid experience. (On a side note, why do companies skip past the number nine? Windows 9, and iPhone 9 where are you?)

Anyway, we hope Microsoft can break the doomed cycle. With wider support for gamers, a great look, and plenty of intuitive features, we can only hope that Windows 11 provides a solid reliable experience for years to come. We have been testing the Windows 11 Home 22000.100 build and are pleasantly surprised. We're certainly excited about the full release, which is expected in late 2021.⏻

PC tech standards are constantly ch
down what to expect from the next g

NEXT-GEN
NIGHTMArES

langing. *Christian Guyton* breaks
generation of computer hardware.

WHEN YOU ACTUALLY EXAMINE how industry standards for connectivity and hardware performance evolve over time, it's surprising how quickly the new normal gets adopted by the masses. USB-A was the dominant form of physical port connectivity across the hardware spectrum for years, but mobile devices ran through a flurry of alternate compact ports before settling on the USB-C connector we know and love.

It's strange to think that in a few short years, the top-tier hardware we feature here in 2021's issues of *Maximum PC* will be heavily outdated, even obsolete. The industry moves fast. When a new gold standard for any aspect of our systems arrives,

it doesn't take long before every new product conforms to the norm. PCIe 4.0 SSDs were hailed as a revolution when they first arrived, exclusive to AMD systems, but Intel has since climbed aboard and heavy hitters such as Samsung have pushed the medium to further heights than the first wave of Gen4 drives could ever reach.

But what comes next? There are plenty of upgrades already in the works, from the next generation of PCIe technology to new DDR5 desktop memory. We're going to analyze where the industry is going next, how soon it's going to get there, and what you can do to make yourself and your hardware as ready as possible.

WE'LL START WITH PCIE, or 'Peripheral Component Interconnect Express'. A bus standard that is commonly found as an interface on motherboards, PCIe connections support a wide variety of components, such as GPUs, SSDs, and Ethernet adapters. The electrical interface provided by the PCIe specification supports other existing standards within the PC hardware industry, including SATA and M.2.

PCIe specifications are developed by the PCI Special Interest Group, a consortium of over 900 tech companies with a vested interest in the technology. The standard is currently in its fourth iteration, although it replaced the original PCI bus standard, which went through five full revisions itself (from PCI 1.0 through to PCI-X 2.0). These standards are constantly evolving; PCIe may birth an entirely new bus format in the not-too-distant future.

For now, we know what to expect from upcoming versions. The specifications for PCIe 5.0 were finalized and released by the PCI-SIG in May 2019, after two years of internal research and development. Preliminary specs for PCIe 6.0 were announced just three weeks later, but this is still in the early testing phase. PCIe 5.0, on the other hand, has already entered production.

In late 2019, Chinese company Jiangsu Huacun Electronic Technology revealed the world's first PCIe 5.0 controller, named the HC9001. Chip manufacturer Rambus announced the completion of a PCIe 5.0 interface in 2020, with others following suit. These technologies are not yet available for use in conventional computer systems, but they aren't far off. AMD stated last year that it expects to have PCIe 5.0 compatibility baked into its 5nm Zen 4 processor architecture, with compatible drives arriving for consumers in 2022.

PCIe 5.0 tech is already poised to take over a number of enterprise sectors, from cloud computing to supercomputers. NAND flash controller producer Silicon Motion isn't quite as optimistic as AMD, saying it expects enterprise-grade SSDs to be available in 2022, with consumer models a little further down the line. Marvell announced earlier this year that it has completed

PCIᴇ VERSION BANDWIDTH

	Link Speed (GT/s)	x16 Bandwidth (GB/s)	Lane Bandwidth (GB/s)
PCIe 3.0	8	32	1
PCIe 4.0	16	63	2
PCIe 5.0	32	126	4

theoretical speeds are even higher, but that would require a 32-lane interface, which are extremely rare and unlikely to become common on consumer products for quite some time.

126GB/s is an incredibly fast transfer rate for an interface that could power desktop graphics cards and drives. PCIe 5.0 M.2 SSDs are likely to become available in the next few years, but will require a processor and motherboard with Gen5 compatibility to perform at its peak. M.2 interfaces use a four-lane connection, which will see the potential maximum transfer speed for SSDs rise from 8GB/s to 16GB/s when Gen5 drives arrive.

What Do You Want?

If you're thinking 'that's much faster than I need!', you're probably right! PCIe 5.0 drives will be monstrously fast, but Gen4 SSDs are already seriously nippy; if you have one of Gigabyte's Aorus 7000s drives, you're looking at read speeds of up to 7GB/s and write speeds of up to 5.5GB/s. That's fast enough to make regular file transfers and load times in games trivially quick.

Upgrading to a Gen5 drive that offers 16GB/s read speeds— or more likely something in the realm of 14GB/s for the first wave of PCIe 5.0 SSDs, as manufacturers perfect the controller hardware—might make a difference, but you won't notice it

If you think 'that's much faster than I need!', you're probably right. PCIe 5.0 drives will be monstrously fast, but Gen4 SSDs are already nippy.

its first PCIe 5.0 SSD controllers under the Bravera product line, which are capable of up to 14GB/s read and 9GB/s write.

Speed Demons

What does PCIe 5.0 actually mean, though? The main benefit here is staggeringly fast speeds; 5.0 offers effectively double the data transfer rate of 4.0, as 4.0 did for 3.0. The link speed is 32GT/s (that's 'giga-transfers per second') compared to PCIe 4.0's 16GT/s, which in practical terms translates to an aggregate bandwidth of a whopping 128GB/s on a standard sixteen-lane interface.

These 'x16' interfaces are found on most modern motherboards for GPUs and other expansion cards, and real-world speeds are likely to be around 126GB/s to account for the overhead penalty in data encoding—this is around 2 percent, as PCIe interfaces from 3.0 onwards use a 128b/130b encoding scheme. The

unless you routinely move massive files around on your PC. Games might load half a second faster, but that's not really a big deal if they already only take two seconds to load.

Power draw on PCIe 5.0 drives is expected to be a lot higher than their Gen4 counterparts, although this is based on current non-consumer products being tested and the thermal performance will likely be refined before we can get our hands on them. Gen5 SSDs are also expected to have superior lifespans compared with previous generations, but again, most SSDs will already last until they're decidedly outdated anyway.

We'd advise holding off on upgrading to PCIe 5.0 unless you regularly need to shift large amounts of data; even then, when it comes to uploading to a network, you're going to be bottlenecked by your internet speed long before your drive starts to struggle. Maybe in several years, games will be big enough to need PCIe 5.0 drives for speedy load times, but that's likely a long way off yet.

Eight-stick kits like this could be a thing of the past with DDR5's increased capacity.

Doubled Data

Moving on from PCIe, let's take a trip down memory lane, literally; we're looking at RAM, specifically DDR5. That's DDR5 SDRAM, the kind used for actual system memory, rather than the memory found in graphics cards (graphical memory is currently in its sixth iteration, GDDR6). A quick note on graphical memory; while GDDR7 or a similar upgrade is bound to arrive eventually, the limits of GDDR6 are still being pushed and we shouldn't expect to see it announced anytime soon.

DDR5 SDRAM—the mouthful that is Double Data Rate Five Synchronous Dynamic Random-Access Memory—is the next generation of desktop RAM. The prevailing industry expectation is that most systems using DDR4 memory will migrate to fifth-generation SDRAM in the future, but how long will that take?

DDR5 memory has been in production for a while now, with members of JEDEC (the Joint Electron Device Engineering Council, a similar conglomerate to the PCI-SIG) announcing work on DDR5 chips as far back as 2017. The standard was delayed from its original planned release in 2018, and was officially released in late 2020 with the launch of the world's first proper DDR5 RAM by Korean semiconductor giant SK Hynix.

Much like PCIe 5.0 products, DDR5 isn't commercially available yet for hardware tinkerers. The same goes for LPDDR5, the low-power memory variant designed for use in laptops, tablets, and smartphones. Although Team Group has released a kit of consumer DDR5 RAM, it's prohibitively expensive and effectively unusable until processor and motherboard manufacturers release compatible boards, hopefully before the end of 2021.

Powered Up

So, what can we hope to see from DDR5 memory? For starters, the memory speed is almost comically fast. The base frequency of DDR5 RAM will be 4,800MHz (that's technically 2,400MT/s, thanks to the doubled data rate), with potential for a ceiling of 8,400MHz or higher, although initial products are expected to launch with a cap of 6,400MHz until upgrades to the specification can be implemented as they were in previous generations.

With comparable latency to DDR3 and DDR4, that puts base memory speed of DDR5 RAM above the vast majority of memory kits currently available for consumers. It is expected that high-end DDR5 kits will run at speeds exceeding even the most aggressively overclocked DDR4 memory. Chinese memory manufacturer Netac teased the development of a 10,000MHz kit this year, which sounds absolutely wild.

Memory capacity per DIMM slot is also being boosted, doubling from 64GB to 128GB. If we see consumer motherboards with the capacity for eight DIMMs as we have for the DDR4 generation—these are rare, with four or two slots being much more common, but mobos such as the Asus ROG Rampage VI Extreme Encore offer eight memory slots for extreme performance—there is potential for monstrous 1,024GB system builds. Imagine how many Chrome tabs we could open with that kind of power...

Is that necessary? Heck no. But the implications are clear; 32GB and 64GB sticks will become more common, raising the bar in particular for compact systems. Being able to easily slap 128GB of memory in an ITX system with only two DIMMs is a huge potential boon, although we've yet to see the software and gaming industries step up to demand such huge amounts of RAM for common computing processes.

One of the most interesting benefits offered by DDR5 is the lower operating voltage of 1.1V, compared with 1.2V in the current DDR4 memory. This results in lower power consumption per DIMM despite the improved performance, which will certainly be a boon for watt-hungry machines.

DDR5 sticks can potentially have voltage regulators built directly into each unit, further boosting data speeds by rapidly adjusting the voltage supplied to the RAM chip on the fly, but this is liable to make manufacturing costs skyrocket. Don't expect to see this feature in standard consumer chips, although we won't be too surprised if kits aimed at overclocking enthusiasts surface.

Caldigit's Element Dock was one of the first USB4-powered hubs to hit the market.

PCIe 4.0 drives felt insanely fast when they first arrived, but they'll soon be overshadowed.

Tech Tricks

The DDR5 standard has several clever inclusions that promise to boost performance compared with previous generations. One of these is a nifty technique called Decision Feedback Equalization. In simple terms, DFE allows for performance improvements by stabilizing data transfers within the DIMM itself and between the memory and CPU.

Looking at DFE in greater detail reveals exactly how it benefits memory performance. With standard SDRAM, signals are sent across the motherboard PCB between the CPU and RAM, with a degree of dielectric loss occurring in the process. This introduces minuscule imperfections in the signals received by the memory, referred to as intermodulation distortion.

DFE circuitry in DDR5 memory introduces a compensating filter that analyzes the distortion and produces a 'counter-signal' that cancels the frequency variation, resulting in a more ideal channel response. Less distortion allows for higher memory I/O speeds, which means greater bandwidth. DFE isn't the only equalization technique used in modern RAM, however, it does represent a significant step forward in performance for DDR5; new memory using DFE should be able to achieve a bandwidth of up to 51.2GB/s per module.

Another pair of features for correcting potential lapses in memory performance are ECC and ECS: that's Error-Correcting Code and Error Check and Scrub, respectively. These are essentially baked-in memory chip functions that search for errors as they occur and rectify them before sending the data to the processor. Some specialized DDR4 kits have previously included ECC tech, but both ECC and ECS are part of the DDR5 memory standard.

ECC-enabled DDR5 RAM will contain separate data lines for the sending of error data to the CPU, allowing additional issues to be corrected by the processor in the nanoseconds of data transit time between the components. This means that DDR5 systems will benefit from improved overall memory stability compared to their DDR4 predecessors.

Leaks Aplenty

What hardware will you need to run DDR5 memory? Well, you won't have any hardware that can run this new memory yet; this isn't a case of simple BIOS firmware updates to accommodate a next-gen CPU. Thanks to a bevy of roadmap leaks (and eventual official announcements), though, we do have a rough idea of when to expect the arrival of DDR5-compatible hardware.

Intel's 12th-generation Alder Lake CPUs will use a hybrid architecture under the new 'Intel 7' process and are expected to launch later this year. The upcoming Sapphire Rapids CPUs (the 7nm Xeon equivalent of Alder Lake, for large-scale business use in servers and supercomputers) will be released later, projected for early 2022. Both will support DDR5 memory, meaning that Intel is almost certainly going to beat AMD to the punch.

Team Red might have a flashy new 5nm architecture planned for Zen 4, but those chips (presumably presented to consumers as sixth-generation Ryzen) aren't expected to release until 2022. Less information is currently available on Zen 4, but it will undoubtedly support DDR5. Some more tantalizing insider info points to Zen 3+ APUs coming sooner with DDR5 compatibility, but all this should be taken with a pinch of salt, as it hasn't been officially confirmed by AMD at the time of writing.

Naturally, motherboard manufacturers will have to keep up with the pace. Current-generation mobos for both Intel and AMD processors—that's Z590 and X570 respectively—don't support DDR5 natively, and news on backward compatibility is patchy. We can probably expect Z690 and X670 chipsets to arrive in late 2021/ early 2022, coinciding with the release of next-gen CPUs. There's not necessarily any guarantee they'll be called that, of course, but it's easy enough to extrapolate naming conventions.

Port Problems

Speaking of naming conventions, we'd like to have some stern words with the bosses of the USB Implementers Forum. After generations of USB 2.0, USB 3.0, USB 3.2, and so on, we've now got a hot new USB standard incoming. Any sane individual might naturally assume this would be called USB 4.0, or perhaps USB 3.3, depending on how much of a technological jump forward the new specification represents.

Micro-USB and Mini-USB are already on the way out; could USB-A be soon to follow?

PAST & PRESENT USB STANDARDS

	Nominal Speed (Gb/s)	Encoding (b)
USB 3.2 Gen 1x1	5	8/10
USB 3.2 Gen 2x1	10	128/132
USB4 Gen 2x2	20	64/66
USB4 Gen 3x2	40	128/132

© GETTY IMAGES/ SUPERSMARIO

RAM SPECIFICATIONS

	Voltage (V)	Bandwidth (GB/s)	Base Data Rate (MT/s)
DDR3	1.3/1.5	17	800-2133
DDR4	1.2	25.6	1600-3200
DDR5	1.1	52.1	3200-6400

Nay, says the USBIF. Chiefs of tech titans such as Microsoft, Apple, Intel, and others have convened, and in their infinite (read: extremely finite) wisdom, have deigned to name this new standard 'USB4'. No .0 here, not even a space or a dash after the B. Nonsense, we say, but what can you do?

The blueprints for the USB4 specification were released over two and a half years ago in February 2019, and USB4 products are already in circulation, albeit in a somewhat limited capacity. Some 11th-gen Intel products have USB4, and USB4-compatible docks are available to purchase. Branding has been broken up into 'USB4 20Gbps' and 'USB4 40Gbps', which are actually USB4 Gen 2x2 and Gen 3x2.

40Gbps throughput means that USB standards have properly caught up to the plucky young upstart that is Thunderbolt, with both Thunderbolt 3 and 4 sharing that same speed. In fact, USB4 represents the closest we've come so far in merging the two interfaces; Thunderbolt 4 mandates the use of USB4 and its full feature set, so any Thunderbolt 4 product will have full USB4 support (although annoyingly the reverse is not true, as we witnessed with USB-C 3.2 and Thunderbolt 3). Intel's somewhat unexpected rollback of royalty fees for using Thunderbolt was a key factor here, opening up the format for other manufacturers.

connection, four for input and four for output. Each lane offers 10Gbps of potential data transfer speed, adding up to a 40Gbps maximum in either direction.

The interesting thing about USB4 is that as a bus specification, it doesn't just provide a simple mechanism for data transfer. Much like Thunderbolt, it can be used to provide tunneling for different protocols, most significantly DisplayPort 1.4 and 2.0. USB4 connectors will be a viable option for display daisy-chaining, with a theoretical maximum transfer performance of 8K at 60fps with HDR10 color enabled.

USB4 cables can achieve this thanks to mono-directional connections, which means that data is only sent in one direction (in this case, from the output system to the monitor), enabling all eight data lanes to be used simultaneously, achieving a transfer speed of 80Gbps, double the USB4 standard. This is only possible with select monitors that are capable of supporting alt-mode DisplayPort 2.0 via USB-C, though.

USB-C ports on new, USB4-compatible devices will have full backward compatibility (albeit at the slower speed afforded by whatever USB-C connector you use). Consumer motherboards for custom PC builds that support USB4 and Thunderbolt 4 are only just starting to seep through the cracks, with Asus's ProArt B550 Creator board being a prime example.

Even though USB4 adoption is still in the early stages, Asus was particularly keen to get ahead of the curve, since USB4 and Thunderbolt 4 are likely to become the prevailing industry standards within a few years.

Power delivery is obviously also a feature here, as seen in previous Thunderbolt interfaces. Using the USB PD format ('PD' stands for 'power delivery'), a USB4 connection can sustain up to 5A of current or 48V of voltage, though not simultaneously. In real-world terms, USB PD is limited to delivering 100W of power at a time, with varying configurations of current and voltage. Minimum values of 1.5A and 5V are applied.

DFE circuitry in DDR5 introduces a filter that analyzes distortion and produces a 'counter-signal' to cancel the frequency variation

Goodnight, Sweet Prince

Yes, you read that right: USB4 is switching over to the same physical connection standards as Thunderbolt 3 and 4, marking the death of the humble USB-A port. It's USB-C from here on; no longer will we struggle with inserting flash drives the wrong way up. No longer will we suffer with non-uniform display cables! Cast off your chains! The revolution is here!

Yes, this will probably lead to some confusion unless the USB and Thunderbolt interfaces are eventually unified; a USB-C port on a laptop right now could be USB4, or Thunderbolt 4, or Thunderbolt 3, or USB 3.2... checking specification sheets for products has always been a must, but this makes it more necessary than ever.

The maximum speed of 40Gbps won't be available on all USB4 products; budget laptops and their ilk are likely to feature the slower 20Gbps ports. This has to do with the available USB lanes for data transfers. USB4 specifies eight lanes in a standard

Whether USB4 could be used for greater power delivery in the future is still up for debate, as thermal limits will quickly become a problem. A USB4 interface could theoretically deliver up to 240W of power, but we aren't likely to see any hardware capable of hitting those figures anytime in the near future.

Getting Set

USB4 is a far more prominent change to the way we build our PCs than the likes of PCIe 5.0 and DDR5. While the last two are more about ensuring your motherboard and processor are up to date, a fully USB4-supported custom build will require an appropriate case I/O as well as matching internal components.

The eventual widespread implementation of USB4 and Thunderbolt 4 will also likely have a knock-on effect on the design of PC peripherals. Despite the format's current dominance, USB-A connectors are liable to die out eventually, mandating

either a system upgrade for better USB-C connectivity or the cumbersome use of numerous adapters. Expect products such as headsets, keyboards, and even monitors to slowly shift to USB-C as a primary wired connection format in the coming years.

Given that pre-built USB4 products are already on the market, it's a safe bet that we'll see support for the new standard in PC components fairly soon. If you're a fan of wireless setups (we love a good Bluetooth mouse), you might not need to worry, but if your I/Os see heavy use we'd advise you to wait until compatible parts are available before doing any major system upgrades.

Blue Teeth

DDR5, PCIe 5.0, and USB4 might be the most pertinent new standards we can expect to see rising to prominence in the near future, but that's not all there is to look forward to. New connectivity standards are always in development; Bluetooth is due an update, with the most recent new format (Bluetooth 5.2) released back in 2019—and that was merely a refinement of the five-year-old Bluetooth 5 standard.

Will we see Bluetooth 6 next, or merely Bluetooth 5.3? Previous trends suggest it will be the former; every preceding generation of Bluetooth technology has stopped short of a .3 iteration, with the Bluetooth Special Interest Group typically pushing ahead with a new standard after one or two updates.

What could the sixth iteration of Bluetooth tech look like, then? We're firmly in speculative territory here, but reduced battery consumption is almost certainly going to be a factor. Bluetooth 5.2 is sometimes referred to as Bluetooth LE, which stands for Low Energy; we've all been there, realizing we left our Bluetooth

on and finding our phone battery is dead. 5.2 introduced a number of power-saving features, which are bound to be improved upon in future versions of the technology.

Beyond LE functions, we can likely also expect improvements to the maximum range and data rate of Bluetooth devices. Bluetooth 5 included functionality that allowed for boosted connection speeds at the cost of range, or a massively improved range (up to four times the standard radius) with a slower data rate. The sixth iteration is expected by speculators to develop this further, with the possibility of Bluetooth-powered one-to-many broadcasts over much longer distances.

Heads Up, DisplayPort

DisplayPort 2.0 is another standard we can expect to see more of soon, with the alt-mode bundled in some cases with USB4 as we mentioned earlier. The standard was finalized in 2019 and represents the first major update to the DP connection standard since March 2016 (the release of DP 1.4).

Normal DisplayPort connectors will remain prevalent for a while, so it's unlikely that you'll need to urgently make any changes to your machines to accommodate the new standard. DP 2.0 offers greater flexibility than any previous version, with support for resolutions above 8K and improved HDR options. Standard DP 2.0 ports will be able to support a blinding 16K resolution at 60Hz, or up to three 8K displays at the same refresh rate.

Is this of interest to the average consumer? Right now, absolutely not. 16K displays are currently huge (far too big for a desktop setup) and only available to corporations via special orders. Sony offers an 18ft tall model, but it will set you back a

Wi-Fi 7 will require new hardware, but isn't due to be finalized until at least 2024

few million bucks. Even the prospect of multiple 8K screens feels outlandish to us; it's not a standard that is representative of the common man's hardware arrangement, so DP 2.0 is reaching for the stars and leaving most of us behind.

Still, as display technology has been refined we've witnessed a steady shift upwards to higher resolutions as the default option. 1080p is currently the most popular resolution for PC gamers, but it hasn't held that crown for that long; plenty of people still play at 720p, while virtually every game has proper support for resolutions up to 4K now. 8K screens could well be the norm within ten years. That said, we'll need the current GPU crisis to settle down before the public can start to reap the benefits of higher resolutions.

Wi-Fi, on the other hand, is almost guaranteed to continue evolving. Specified as the IEEE 802.11 standard, a subset of IEEE 802, WLAN standards are diverse in purpose and functionality. Most 802.11 iterations are used for specific functions beyond traditional Wi-Fi as we might use it in our homes; for example, 802.11ah defines a wireless network running at sub-1GHz bands, ideal for wide-range, low-power uses such as outdoor Wi-Fi for traffic control.

The Wi-Fi we know and love is 802.11ac, or 802.11ax for those with more recent hardware. The 'ax' iteration is more commonly known as 'Wi-Fi 6' and represents the current apex of consumer WLAN options, thanks to better range and superior performance in signal-dense environments.

The Wi-Fi we know and love is 802.11ac, or 802.11ax for more recent hardware. The 'ax' iteration is more commonly known as 'Wi-Fi 6'

High Definition

HDMI acts as a sister format of sorts to DisplayPort, found more commonly on TVs than computer monitors. Unlike DisplayPort, there isn't a new standard of HDMI incoming; the most recent version is HDMI 2.1, officially released in 2017. The space between HDMI versions is typically greater than DP iterations, and HDMI 2.1 is already capable of supporting up to 10K resolution at 120Hz.

It's entirely possible that homogenization of display connectors could occur over the next few decades, though only if the wider industry is willing to play nicely. The USB-C connector is already commonplace and could represent a shared ground for data transfers and display support, although this would ultimately mean the eradication of HDMI as a display interface standard.

What other evolving standards might we see in the future? The only key ground we haven't covered here is internet connectivity; Ethernet and Wi-Fi. The IEEE 802 family of standards (maintained by the Institute of Electrical and Electronics Engineers) has remained largely unchanged for years, presiding over LAN and PAN connections, amongst others. Physical cables for Ethernet connections have been improved over time, but the classic Ethernet port is unlikely to be leaving us soon.

Anticipating Wi-Fi 7 is a tricky business. Wi-Fi 6E was only announced a few months ago, offering 6GHz band operation compared to Wi-Fi 6's 2.4GHz and 5GHz options. There are a few different IEEE 802.11 standards currently in development, but the pertinent one is 802.11be. Also referred to as EHT (Extremely High Throughput), this standard will focus primarily on building upon the 802.11ax format with massively superior maximum speeds and improved signal stability.

Projections put the potential maximum throughput of 802.11be at a seriously nippy 30Gbps, triple that of Wi-Fi 6. The most common expected use is video streaming, which has increased massively in recent years. Using Wi-Fi 7 will demand compatible hardware, so be ready to buy a new router and motherboard. However, the IEEE doesn't expect to finalize the standard until 2024 at the earliest, so we'd consider this a low priority right now.

The world of PC hardware and software never stops marching on. We've made some predictions here that might never ring true; for all we know, a new standard could arrive on the scene next week and disrupt the status quo. Will it be a new bus interface? A new type of graphical memory? A port for video output? We couldn't say, but we can't wait to find out what comes next.⏻

SUBSCRIBE TO
MAXIMUM**PC**

GREAT REASONS TO SUBSCRIBE

- Save over 74% when you subscribe.
  - Free *PC Gamer Annual 2021* digital bookazine.
    - Never miss an issue.
    - Get every issue delivered direct to your door.

TO ORDER GO TO
www.magazinesdirect.com/mpc/B1JMAG

EVENT ON THE
HORIZON

Ian Evenden explains how the Sun could one day take out our electronics and communications technology in mere moments

EARTH, SEPTEMBER 1ST, 1859. Colors flash through the night sky above New England, gold miners in the Rocky Mountains are woken by the brightness of the Northern Lights, visible as far south as the Caribbean. Telegraph operators across the world receive electric shocks from their equipment, which continues to operate, despite being disconnected from the power supply.

The Sun, August 31st, 1859. A complex system of magnetic field lines suddenly twists, releasing a large quantity of plasma into space. This takes 17 hours to cross the 93 million miles to the Earth, which is at just the right place in its orbit to be hit by what today we'd call a coronal mass ejection.

The largest geomagnetic storm on record, the Carrington Event caused widespread electrical disruption and power blackouts in an electrical grid that was primitive compared to today's complex system.

Should it happen again, the consequences could be catastrophic. A 2013 research project from Lloyds of London and Atmospheric and Environmental Research in the United States estimated the cost to the US alone could be $2.6 trillion.

At the peak of its activity, the Sun belches out as many as three coronal mass ejections every day. One only just missed us in 2012, and if it struck today, the damage would be incalculable.

WHAT'S IT ALL ABOUT?

In their day-to-day lives, our PCs and other electrical equipment are unlikely to come into contact with charged particles, but every now and then, the Sun reaches out to touch us. Protected in the Earth's magnetic bubble, we don't often notice the effects of the solar wind unless we live far enough north (or south, hello readers in New Zealand) to see the aurora. Our Sun is, compared to other places in the Universe, a relatively placid, middle-aged star, but occasionally it can surprise us.

The Sun operates on an 11-year cycle. In 1859, it was approaching the middle of this cycle, the time of greatest activity. Astronomers, equipped with ever-improving telescopes, were starting to take more of an interest in the Sun around this time, and the first observation of a solar flare was made on September 1st that year by the astronomers Richard Carrington (for whom the solar storm is named) and (independently) Richard Hodgson, both based in southern England.

That flare, which Carrington observed by projecting the output of his telescope onto a screen through a broad-band filter (remember, never look at the Sun with the naked eye, or with any kind of magnifying equipment, or indeed with anything other than an approved solar filter) turned out to be enormous, a white light flare of extraordinary intensity.

Solar flares are associated with coronal mass ejections, and both are common when sunspots are on display, as these temporary, dark patches are signifiers of magnetic activity on the star's surface. As the 11-year cycle goes on, the sunspot count moves from none, sometimes for hundreds of days at a time, to anything up to several hundred at once.

OF SUNSPOTS AND CMES

Nobody was counting sunspots in 1859, though they were known to Chinese astronomers back in antiquity, and were mentioned by the Ancient Greeks. They were first drawn (that we know of) by an English monk in 1128. It took until 1610 to get a telescope on them, which is fair enough as the instrument was only patented in 1608, but it wasn't until the early 1800s that the astronomer William Herschel was able to associate sunspots with varying levels of solar activity. His hypothesis that an absence of sunspots led to higher wheat prices on the market, was widely ridiculed at the time.

While the precise nature of sunspots is still a matter of research and debate, it seems like he was right. Solar minima lead to cooler years, on average, which

Above: **These 366 images of the Sun, taken by the ESA's Proba-2 satellite, represent every day during 2020. Activity comes and goes, including two partial eclipses.**

Left: **The STEREO probes are an identical pair of spacecraft designed to observe the Sun. They orbit ahead of and behind the Earth to get a stereoscopic view.**

would have made wheat harvest smaller, pushing prices up. Fossil records suggest this cycle has been stable (between nine and 14 years, for an average of 11) for 700 million years. We're currently in an active, warm period, similar to that around the year 1000. There have been extended cool, low activity periods in the past too, most recently the 'little ice age' that ran from the 16th to 19th centuries.

Sunspots appear to be the visible effect of magnetic flux tubes in the Sun's convective zone—an unstable layer just below the star's surface that's churning with convection currents as heat and other forms of energy are sent out by the fusion reactor running at its core. They are twisted and wound by differential rotation—different areas and depths of the Sun don't all rotate at the same speed.

So what happens when the irresistible force of the Sun's plasma ejection hits the immovable object of the Earth's magnetic field? "These large eruptions release large amounts of hot material called plasma," says Dr Ravindra Desai of Imperial College London, who previously worked at the Los Alamos National Laboratory, New Mexico, and NASA Goddard Space Flight Centre, Maryland. Among other interests, Desai attempts to forecast the Sun's activity, a phenomenon known as space weather.

"When the plasma strikes Earth, it's deflected by the planet's magnetic field," he says. "The Earth protects us from the majority of the onslaught, but the plasma flows along those magnetic field lines and causes the Aurora Borealis which is a wonderful thing to see, but it also causes [electrical] current systems in the ionosphere and the ground. They aren't dangerous to human beings, but can knock out power stations and satellites."

© ESA, NASA

Above: **This solar flare was photographed by NASA's Solar Dynamics Observatory in 2017. Here, you can clearly see the coils of the Sun's magnetic field lines.**

This happened in 1989, when a billion-ton cloud of plasma was ejected from a powerful explosion on the Sun at a million miles per hour. Two days later, it struck the Earth, sending aurora as far south as Cuba and causing the entire Canadian province of Quebec to lose power for 12 hours after its transformers blew.

Electrical grids across the United States also saw drops in their power output, with over 200 incidents reported. Luckily, the US had enough power in reserve to keep the lights on. Satellites tumbled out of control, and the orbiting Space Shuttle Discovery developed a mysterious fault in one of its hydrogen tanks that disappeared at the same time the storm abated.

While the 2012 event missed the Earth by approximately nine days, it instead hit the STEREO A spacecraft, part of a pair of sun-observing probes that orbit both ahead of and behind the Earth to provide a stereoscopic image of the Sun. STEREO A (the probe ahead of the Earth) was largely undamaged by the plasma due to not being within a magnetic field at the time, but the failure of STEREO B led to the ending of the mission in 2018.

The 2012 event is thought to have been as powerful as in 1859, and it looks like Earth had a narrow escape. "If it had hit, we would still be picking up the pieces," said Professor Daniel Baker, director of the Laboratory for Atmospheric and Space Physics at the University of Colorado in a NASA statement in 2014. "In my view, the July 2012 storm was at least as strong as the 1859 Carrington event. The only difference is, it missed. I am convinced that Earth and its inhabitants were incredibly fortunate that the 2012 eruption happened when it did. If it had occurred one week earlier, Earth would have been in the line of fire."

CHANCES OF IT HAPPENING AGAIN

Also writing in 2014 was the physicist Pete Riley, of Predictive Science Inc. He analyzed records of solar storms going back over 50 years, and by extrapolating the frequency of ordinary storms to the extreme ones, he calculated the odds that a Carrington-class storm would hit Earth in the next ten years. His answer: 12 percent, or one in 8.3. We may have three years left on this prediction, but the threat doesn't go away after that.

As it happens, we're now entering a new solar cycle, the 25th since records began. This means that solar activity has been at a minimum for a while, "but in the last few months we've already started to see the Sun become more active," says Desai. Cycle 24, which ended in 2019, was a quiet

SPACE WEATHER FORECAST

"At the moment we have telescopes looking at the Sun," says Desai, "they see stuff happening on the Sun's surface, and as soon as a CME occurs we can start running our models to predict what time the CME will arrive, and when it does, what will its properties be—how big is its magnetic field, and how much plasma is contained inside of it. Typically, CMEs can take between three and five days to reach us, but a severe one could reach us sooner."

Indeed, NASA figures put the shortest time for a CME to cross the 93 million miles between us and the Sun as 15 hours, for a mind-boggling speed of over six million miles per hour.

Desai's team has discovered that, if CMEs come at us in pairs, then the second storm will be more violent as a result. While investigating the 23 July 2012 event, they discovered a previous CME on 19 July. The four-day gap between the storms was long enough for the solar wind to recover from the stress, but by using computer modeling to place the second CME closer to the first, the Imperial team discovered that it became stronger, and faster.

If a CME is moving faster than the background solar wind (a million mph), then ahead of the plasma cloud travels an interplanetary shockwave that can arrive at Earth with just an hour's warning. While this shockwave transfers its energy mainly through electromagnetic fields rather than particles, it can push particles ahead of it, increasing the chances of a radiation storm here on Earth.

© GETTY IMAGES, NASA

cycle, yet included the enormous storm of 2012. Predictions have been made for Cycle 25, but anticipating the actions of a star is tricky, and different methods produce different results. "It's already releasing coronal mass ejections," says Desai, "so the chances of a severe solar storm are getting higher."

Those, as every XCOM 2 player knows, are the kind of odds that do come up. "The effects could be pretty bad—it could cost the world's economy trillions of dollars," says Desai. "In terms of actual damage, it could knock out satellites, which would have an effect on our communication systems, anything from financial systems to those that open train doors. Also, electrical current systems in the atmosphere and on the ground can knock out large transformers, which is what happened to Quebec in 1989. Luckily, they were able to get things back online, but these transformers take many months to build. So if a huge solar storm happens, and knocks out these transformers, a worst-case scenario could be that it takes many months, if not years, to get our grid capacity back up to where it was before."

COSMIC RAY PROCESSOR LOCKUP

Energetic particles don't have to come from our Sun, though. Cosmic rays, high-energy protons, and atomic nuclei zipping through space at almost the speed of light can come from as far away as other galaxies, possibly from supernovae. The most powerful are carrying an energy level 40 million times that produced in the Large Hadron Collider, described as like being hit by a baseball traveling at 56mph. Most, however, don't have this kind of energy, but when they hit our atmosphere they still cause a shower of secondary particles, some of which will reach the surface and affect our electronics.

A study by IBM in the 1990s suggested there was one cosmic-ray-related error per 256MB of RAM per month. In a machine with 64GB of RAM, that's 250 cosmic ray errors a month or 8.3 a day. In 2008, Intel patented an on-chip cosmic ray detector that repeated commands it decided had been affected by charged particle interference. Wherever they come from, energetic particles cause electrical currents to move through our atmosphere, and also in the tiny electronics of our CPUs and other devices, providing power where it's not wanted, in the same way the telegraph operators were unable to switch off their equipment in 1859. Whether it's burning out your logic gates or frying the nearest substation, CMEs are bad news for electronics.

What, though, happens if you're not protected by the Earth's magnetic field? Astronauts aboard the International Space Station have to take shelter in the more heavily protected (often Soviet-built) parts of the station during a solar storm, and they're only 250 miles above the Earth's surface. If you're making a probe that will travel further afield, then you need to do something to protect it from radiation. Something like the European Space Agency's Sun-observing Solar Orbiter, for example.

"There is no plan to avoid coronal mass ejections," says Piers Jiggens of the ESA's Space Environments and Effects section. "A major high-level goal of the Solar Orbiter mission is to observe eruptions on the Sun and then to correlate those observations with in-situ measurements.

Below: Aurora over Iceland in 2012 after a CME. As charged particles from the Sun collide with our atmosphere, different colors are generated—green for collisions with oxygen, and red for nitrogen.

From an instrument standpoint, we require that they can observe the [entire] range of plasma and radiation thrown at them, which at the closest passage of Solar Orbiter can indeed be 10-20 times more intense than what we might see close to the Earth. It's more interesting for us to be able to predict them, in case it could be used in operational planning, so we get back the most interesting data.

Plasma can cause charging on the spacecraft's surfaces, so it's important that we use conductive surface materials, paints, and coatings to dissipate this charge and avoid potentially damaging discharges. The higher-energy radiation

An artist's impression of the ESA's Solar Orbiter probe at the Sun.

can penetrate the spacecraft shielding and deposit energy in dielectrics and other spacecraft materials, affecting their electrical and material properties. We need to take this into account when designing the electronics for spacecraft because how they can perform at the end of their lives can be different from at the beginning. This often means operating components in safe zones at lower voltages—a process known as de-rating. The short-term impact of solar radiation storms includes upsets in electronics and memory, which can manifest as non-destructive events requiring correction, or destructive events requiring mitigation."

© ESA

If you're designing a computer to run in space, then you need to build it for the environment. Meet the Rockwell DF-224, 110lbs of silicon that measures 18x18x12 inches and runs at 1.25MHz. This was the original computer fitted to the Hubble Space Telescope and was considered for use on the Shuttle too. It's a highly redundant design, both to mitigate the effects of radiation and because maintenance in space is hard. So there are three CPUs, one active and two backups. Six memory units, of which four could be powered up at once, three I/O controllers, and six independent power converters. Aboard Hubble, the first

servicing mission upgraded the DF-224 with a 16MHz 80386 as a co-processor, then the third servicing mission swapped the lot for a 25MHz 80486. Hubble was recently restarted following a switch to its backup Power Control Unit, showing the resilience of this 90s-era technology.

SPACEPROOF CPUS

Making a processor that can be used in space—as opposed to the relatively balmy environment of the International Space Station, where Linux-based ThinkPads are the norm—means more than just redundant design. Take the BAE Systems

RAD750, the CPU in various space telescopes, the STEREO probes, and the Perseverance rover on Mars. It's based on an IBM Power design known by Apple as the G3, is manufactured on either a 250nm or 150nm process, and with its attendant motherboard is good for temperatures of -67°F and a radiation exposure of 1,000 grays (five would kill a human). The RAD750, a chip that graphic designers gave up in the 2000s, is state of the art for space CPUs, and costs $200,000.

You'd get better performance for your space-buck if you used off-the-shelf components, but unfortunately, this has been tried. A Russian attempt to explore

the Martian moon of Phobos, known as Fobos-Grunt, prioritized budget over spaceworthiness, and 62 percent of its chips were not qualified for going into space. Some, the SRAM, had been tested in a particle accelerator and were found to be extremely vulnerable, latching-up (a short circuit where a voltage spike in one transistor can spread to others nearby) at the minimum levels of particles it was exposed to. Still, the Russian space agency stuck it on top of a rocket and fired it at Mars in November 2011.

Explanations differ, but one story, sourced from a Russian military newspaper, is that a particle passed through the SRAM preventing the probe from firing its engines to get out of the low Earth orbit provided by the Zenit-2M rocket. The dual processors in the onboard computer initiated a reboot, then went into safe mode to await instructions. These never came, because the antennas the probe used to listen to ground control were extended in the cruise stage after it left Earth orbit. No antennas meant no engines, and no engines meant no antennas. The probe, worth 2.4 billion rubles (approx $81 million), fell into the Pacific ocean in January 2012.

GOING NUCLEAR

There's another way that our electronics could end up on the wrong end of the electromagnetic force, and that's thanks to a nuclear explosion. Like a coronal mass ejection, the EMP generated by a nuke is a short-term phenomenon rather than the constant bombardment of radiation in space. A nuclear device detonated at the right altitude could bathe the entire continental US in EMP, having a "catastrophic" effect on communications and the grid infrastructure, according to a 2008 report from the EMP Commission.

The group was formed in 2001 to "assess the nature and magnitude of potential high-altitude EMP threats to the US from all potentially hostile states," and its report (https://bit.ly/MPC_boom) makes for terrifying reading. It discusses how everything from water and food supplies to transport and fuel distribution would be affected by a burst of charged particles from the upper atmosphere. A high airburst would inflict maximum damage—explode close to the ground and it loses much of its energy into the Earth.

As with CMEs, ionized particles passing through the chip cause currents that can lead to transistors randomly opening, or latch-ups. These phenomena cause bit-flips, where a 1 or 0 becomes a 0 or 1, changing the content of memory or the

data being processed. Using ECC memory is one way of dealing with this, and the errors can often be cleared with a power cycle, but enough voltage can cause a transistor to burn out, and long-term exposure causes gradual degradation of a processor or memory chip as electron-pair holes build up in the gate insulation layers. Eventually, the charge built up by charged particle bombardment can keep transistors to stay open or closed. The smaller your production node, the more likely you are to suffer the effects of charged particle damage, so as our tech advances, we make it more likely to fail in the event Earth is hit by a CME or EMP.

Above: **This coronal mass ejection was recorded by the Solar Dynamics Observatory in February 2015. A substantial part shot off into space.**

Left: **BAE Systems' RAD750 CPU is based on the same IBM architecture as Apple's G3.**

Chips designed to mitigate these problems are often built on an insulating substrate rather than the normal semiconductor wafer, such as silicon-on-sapphire, which uses aluminum dioxide's insulation properties to prevent stray currents doing an end-run through the substrate and leaking into more of the chip. Space-grade sapphire chips can withstand up to 3000 gray, and eliminate latch-ups. It's also possible to build in redundancy at the bit level, with each logical bit replaced with three.

Another method is to have three microprocessor boards instead of one, all doing the same calculations and comparing answers. Any minority answers are recalculated, and a board that consistently reports a wrong answer is shut down or restarted. Restarts can be carried out on a timer to clear bit-flips, while a simple radiation shield—a slab of lead or other material in sufficient amounts—prevents particles from

Above: **Magnetic loops known as 'flux ropes' rise up from the Sun in this image from the Solar Dynamics Observatory.**

reaching the silicon. When launching into space, however, every fraction of an ounce counts, so heat shields are made with clever design, keeping as much of the body of the spacecraft as possible between the CPU and the source of radiation.

Keeping a spare laptop under a lead cover may be the only way to guarantee your tech survives Earth being hit by a CME of the scale seen in 1859. How useful that will be in a world devoid of an electrical grid or functioning communications is another matter, but at least you'll be able to play *Minesweeper* until the battery runs out. Older tech, such as ham radios, seem more likely to survive due to their relatively large transistors being less likely to burn out in a voltage spike, but there's a catch: HF (shortwave) signals bounced off the ionosphere for over-the-horizon communications can be disrupted by solar activity such as radio bursts or the X-rays given off by solar flares.

The Sun always seems to be one step ahead of us. At least, with satellites monitoring the Sun and space weather simulations being run on supercomputers, we should get some warning this time. ⏻

CALL A QUANTUM MECHANIC

Quantum computers are notoriously unable to deal with any kind of interference from outside their enclosed systems, often cooled to just above absolute zero for their quantum bits, or qubits, to operate properly. And by properly we mean in a way that makes no sense according to the rules by which we usually understand the universe.

Instead of being limited to the binary digits of 1 or 0, quantum computers are able to reach a 'superposition' of both at the same time, thanks to the peculiar properties of the subatomic particles they're built upon.

The differences in voltage that maintain these states are tiny too, which means anything that can artificially raise or lower a voltage can upset the system.

This is why they're often kept so cold, so that heat from the outside can't seep in and mess up the calculations.

As it stands, quantum computers generally run for a matter of microseconds before something causes them to collapse, although we're getting better at keeping them stable as we add qubits and ramp up the calculating power.

A recent paper published by MIT professors suggests cosmic rays and radiation from the Sun could be a greater threat than anything else, limiting future quantum computers to four milliseconds or so of operating before they decohere.

An energized particle hitting a quantum computer could cause a shockwave effect within a sufficiently complex processor, causing the qubits to fall over like dominos.

The solution? Future quantum computers may have to be built under layers of lead shielding, or underground, to protect them from the natural particle showers.

PLUG & PLAY
PI PROJECTS

Jonni Bidwell explains how to make the most of the versatile Raspberry Pi device.

WE LOVE THE Raspberry Pi here at *Maximum PC*. And so do our colleagues at tomshardware.com, home to *Pi User*'s Les Pounder. So with their help, we've selected the finest Pi projects for you to enjoy. From retro gaming to home security, or even artificially intelligent object classification, there's something here to inspire everyone.

The Raspberry Pi is the perfect device for learning Linux. Its official operating system is based on Debian Linux and there's even a build of Ubuntu for it. The Pi 4 is powerful enough that it can happily replace your desktop, and small enough to hide behind your monitor. Then there's the Pi Zero, which is perfect for budding Internet of Things enthusiasts. And now there's the Pi 400, where the keyboard is the computer.

Of course, there's no need to stop at these projects. We want you to boldly go exploring new computing frontiers, and we hope this offering inspires just that. With a bit of practice and experience, you'll be soldering up hardware or controlling robots—the only limit is your own imagination.

Get your kit ready

Prepare your USB cables, SD cards and a cup of tea. It's time to start looking at some Raspberry Pi projects...

WITH MORE THAN 34 million units sold, the Raspberry Pi is not only one of the world's most popular computers, it's also one of the most important. Designed to help kids learn about technology, this inexpensive, single-board system is the leading choice for makers, developers, and hobbyists who want to do everything from building industrial robots to setting up retro arcade machines. Whether you're aged eight or 80, if you love technology, the Raspberry Pi is for you. And with models costing from £5 to £65, it will suit any budget.

Whatever you do with your Raspberry Pi, you will need an SD card with an operating system on it. The official Raspberry Pi OS (formerly Raspbian) is a popular choice and a great start for many projects. It's available in Desktop and Lite flavors at www.raspberrypi.org/software/operating-systems. There's an official Raspberry Pi Imager utility that you can use to write this (or a selection of other OSes) to an SD card, or you can use the NOOBS tool. We favor Balena Etcher, which has the advantage of being able to handle compressed images.

TROUBLESHOOTING
In this feature, we'll use a few different OS images, so we've made a handy guide to writing them below. If you run into difficulties, it's worth remembering that a faulty SD card might be to blame, so it's worth trying a different one if things go wrong.

Cheap power supplies are another cause of errors. For older model Pis, these are less of a problem, but the Pi 3 and above really do need a

good 5.2V or they'll be subject to CPU throttling and instability. The Pi 4 has switched its power input from micro-USB to USB-C and, unfortunately, this new port isn't USB-C compliant. So if you use a fancy smart charger, it probably won't power the device, and might possibly damage it.

You won't run into these issues if you use the official charger, so do that. If you're opposed to that, we're told that cheaper USB-C cables (ones that lack the smart charge chip) work fine. Cheaper chargers should work too, but may not be so reliable.

We've handpicked some fun, useful, and interesting projects to showcase the tremendous scope of things that can be achieved with the humble Pi. From retro gaming to machine learning and Bitcoin, the possibilities are truly endless...

The latest incarnation, the Pi 400, embodies a Pi 4 computer (pictured) within a keyboard.

WRITING SD CARDS WITH ETCHER

1. DOWNLOAD ETCHER
You'll find Linux, macOS and Windows builds at https://etcher.io. Linux ships as a zipped AppImage, so once unzipped, run `chmod 755` on the file (or select Properties→ Permissions and check the box). Double-click the AppImage to run it (or do so from the CL).

2. PREPARE THE TARGET
Insert your SD card, select the Flash from file option and navigate to your downloaded image. Click Select target and a menu will show available devices (hard drives are hidden to prevent overwriting). Select your SD card reader, and check there's nothing important on it.

3. WRITE THE IMAGE
Hit Flash, enter your password and wait for the progress bar. You'll be offered the chance to write the same image again (eg. if you're preparing a fleet of Pi clones). Otherwise, your SD card is ready. Move it from your PC to your Pi, power on and let the adventures begin!

Image recognition

Get into machine learning and train your Pi to recognise and classify other Pis, without having to write a line of code.

WE'RE GOING TO TRAIN our Raspberry Pi to identify other Raspberry Pis (or other objects) with machine learning (ML). Why is this important? An example of an industrial application for this type of ML is identifying defects in circuit boards. As circuit boards exit the assembly line, a machine can be trained to identify a defective circuit board for troubleshooting by a human.

Other neat applications of machine learning and artificial intelligence could include facial recognition and face mask identification. For these types of projects, you can store the training images locally on the Raspberry Pi, however, the training process may take longer if performed on the Pi. So, for this tutorial, we'll use a web platform called Edge Impulse, one advantage of which is the ease of uploading training images. This can be done from a smartphone, without having to involve an app.

We'll use BalenaCloudOS instead of the standard Raspberry Pi OS, since the folks at Balena have pre-built an API call to Edge Impulse. Some facial recognition and face mask identification tutorials also require tedious command line package installs and Python code. This project eliminates all terminal commands and uses an intuitive GUI instead.

TRAINING ON THE EDGE
Go to https://edgeimpulse.com and create a free account (or login), from a browser window on your desktop or laptop. Select Data Acquisition from the menu bar on the left. You can either choose to upload photos from your desktop or scan a QR code with your smartphone and take photos. In this tutorial, we'll opt for taking photos with our smartphone. Select Show QR code and a QR code should pop up on your screen. Scan it and select Open in browser and you'll be taken to a data collection site. You won't need to download an app to collect images.

Accept permissions on your smartphone and tap Collecting images? in your phone's browser screen. Tap Label and enter a tag for the object that you take photos of. Take 30-50 photos at various angles. Some photos will be used for training and other photos will be used for testing the model. Edge Impulse automatically splits photos between training and testing. Repeat the process of Entering a label for the next object and taking 30-50 photos per object until you have at least three objects. We recommend three to five identified objects for your initial model. You'll have a chance to re-train the model with more photos and/or objects later on.

From the Data Acquisition tab in the Edge Impulse browser window, you should now see the total number of photos taken (or uploaded) and the number of labels (type of objects) you have classified. You can click any of the collected data samples to view the uploaded photo.

IMPULSE DESIGN
Click Create impulse from Impulse design in the left column menu. Click Add a processing block and select Image to add an image to the second column from the left. Click Add a learning block and select Transfer Learning. Click the Save Impulse button on the far right. Click Image under Impulse design in the left menu column. Select Generate features to the right of Parameters near the top of the page. Click the Generate features button in the lower part of the Training set box. This could take five to 10 minutes (or longer) depending on how many images you have uploaded.

Select Transfer learning within Impulse design, set your Training settings (keep the defaults, check Data augmentation box), and click Start training. This step will take five minutes or more, depending on the amount of data. After running the training algorithm, you can view the predicted accuracy of the model. For example, in this model, the algorithm can only identify a Raspberry Pi 3 correctly 64.3 percent of the time and will misidentify a Pi 3 as a Pi Zero 28.6 percent of the time.

Select Model testing in the left column menu. Click the top check box to select all and press Classify selected to test your data. The output of

As well as collecting images from your mobile, Edge Impulse can connect to lots of data sources.

© TOM'S HARDWARE
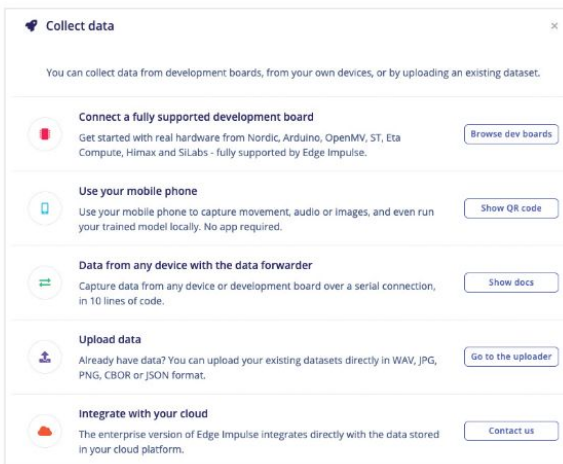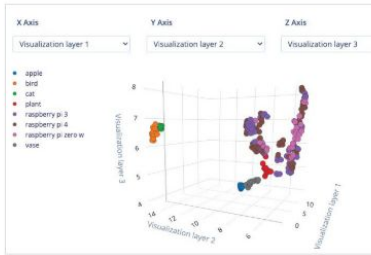
Gather datapoints (images), and Edge Impulse can explore and graph the features of the data.

this action will be a percentage accuracy of your model. If the level of accuracy is low, we suggest going back to the Data Acquisition step and adding more images or removing a set of images. Select Deployment in the left menu column and select WebAssembly for your library. Scroll down (the Quantized option should be selected by default) and click the Build button. This step may also take three minutes or more, depending on the amount of data.

SETTING UP BALENACLOUD
Instead of the standard Raspberry Pi OS, we'll flash BalenaCloudOS to our microSD card. The BalenaCloudOS is pre-built with an API interface to Edge Impulse and eliminates the need to attach a monitor, mouse, and keyboard to our Raspberry Pi.

Create a free BalenaCloud account at https://dashboard.balena-cloud.com/signup and then go to https://dashboard.balena-cloud.com/deploy to open the Create and Deploy page and create a balena-cam-tinyml application. Click Deploy to Application. After creating your application, you'll land on the Devices page. Don't create a device yet!

In Balena Cloud, select Service Variables and add two variables. First, add to the service edgeimpulse-inference a variable named EI_API_KEY and in the Value field paste the API key from the Keys section of the Edge Impulse Dashboard. Add a second variable to the service named EI_

PROJECT_ID and paste the Project ID value from the Dashboard. Select Devices from the left column menu in BalenaCloud, and click Add device. Select your Device type (Pi 4, Pi 400, or Pi 3).

Select the radio button for Development. If using Wi-Fi, select the radio button for Wi-Fi+Ethernet and enter your credentials. Download your customized BalenaOS image and write it to an SD card (using our guide to Balena Etcher on the first page)

CONNECT THE HARDWARE
Remove the microSD card from your computer and insert it into your Raspberry Pi. Attach your webcam or Pi Camera and then power up your Pi. Allow 15 to 30 minutes for your Pi to boot up and BalenaOS to update. You can check the status of your Pi Balena Cloud OS in the BalenaCloud dashboard.

Identify your internal IP address from your BalenaCloud dashboard device. Enter this IP address in a new browser Tab or Window. Place an object in front of the camera. You should start seeing a probability rating for your object in your browser window (with your internal IP address). Try various objects that you entered into the model and perhaps even objects you didn't use to train the model.

We can forgive the machine for thinking this Pi 3 was likely to be a Pi 4. At least it knew it wasn't a plant.

REFINING THE MODEL

If you find that the identification isn't accurate, first check your model's accuracy for that item in the Edge Impulse Model Testing tab. You can add more photos by following the Data Acquisition steps and then selecting Retrain model in Edge Impulse. You can also add more items by labeling and uploading in Data Acquisition and retraining the model. After each retraining, check for accuracy and then redeploy by running WebAssembly within Deployment.

If you want to go further with machine learning, check out Pimoroni's new BrainCraft HAT for the Pi 4. It features a 240×240 TFT IPS display for inference output, slots for a camera connector cable, a five-way joystick, two microphones, audio outputs, and much more. Most importantly, it has a controllable fan to keep the thing cool during all the never-ending TensorFlow computations.

Pi-nky and the BrainCraft HAT will surely aid all your world domination plans.

Multi-room Pi audio

Get your Raspberry Pi playing music around the house.

MULTI-ROOM AUDIO SYSTEMS can cost a pretty penny, but who needs to buy an expensive set of wireless speakers when you can use some Raspberry Pis and any 3.5mm wired speakers to achieve the same effect? We'll use our Pi 4s and speakers/receivers (with 3.5mm audio inputs) to play music from a streaming service perfectly in sync via Bluetooth from our phone, tablet, or computer. We'll show you how to do this with a single speaker and Raspberry Pi and then replicate it in other rooms. This is an excellent project for repurposing old receivers. If you have a non-Bluetooth receiver with a 3.5mm jack input, you can connect your speakers to your receiver and your receiver to your Raspberry Pi.

Instead of Raspberry Pi OS, we'll use BalenaSoundOS to make the Bluetooth connection visible (to our other devices) and eliminate the need for a monitor, keyboard, and mouse for our Pi.

Create a free BalenaCloud account at https://dashboard.balena-cloud.com/signup and log in. Open the Deploy a Balena-sound application page at https://sound.balenalabs.io/docs/getting-started. You must already be logged into your Balena account for that link to work. Click Deploy to Application and select Devices from the left column menu in your Balena-Sound Dashboard. Click Add device and choose Raspberry Pi 4, Raspberry Pi 400 or Raspberry Pi 3. Select the radio button for Development. If using Wi-Fi, select the radio button for Wi-Fi+Ethernet and enter your Wi-Fi credentials. Click Download balenaOS and a zip file will start downloading. Write this to an SD card, using Balena Etcher, or otherwise.

CONNECT AND UPDATE

Remove the microSD card from your computer and insert it into your Raspberry Pi. Connect the speaker to your Pi via the 3.5mm audio cable. Power up your Pi. Allow 15 minutes for your Pi to boot up and BalenaOS to update—only the initial boot requires this long update. You can check the status of your Pi Balena Sound OS in your BalenaCloud dashboard. If you're using a Raspberry Pi 400, connect a USB speaker instead of the 3.5mm audio jack. Wait for your Raspberry Pi BalenaOS update to finish and confirm that all systems are running.

From your smartphone, tablet or computer, navigate to your Bluetooth settings and look for BalenaOS XXXX and pair it to your device. Connect the device as the sound output for your smartphone, tablet, or computer. Go to your streaming service and play a song. The music should stream via Bluetooth to your Raspberry Pi and play from the attached speaker.
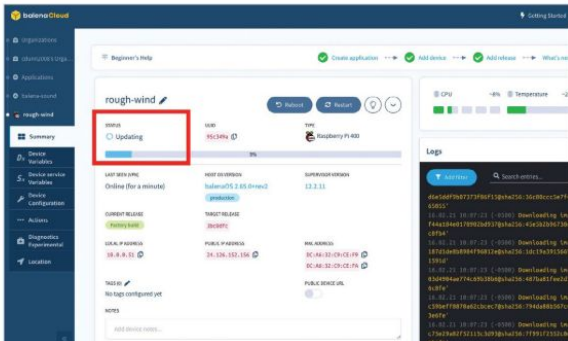
Your musical Pi can be monitored through the Balena Cloud dashboard.

MULTI-DEVICE/MULTI-ROOM SYNC

For multi-device sync, repeat the setup stages we described above for each extra device. You can flash the same image from earlier if you're using the same model of Raspberry Pi. With the same Bluetooth device connected as before, music can now stream to all Raspberry Pis and their connected speakers. You will only need to connect one Pi via Bluetooth to your phone, tablet, or computer to stream to all Raspberry Pi Balena Sound speakers.

That's it! You can now create multi-room and multi-device sound experiences in your home with your Raspberry Pi devices, while DJing from your smartphone, tablet, or computer like a pro.

Audiophiles probably won't care much for Bluetooth audio, nor the headphone output on the Raspberry Pi. Fortunately, there are a variety of DAC (digital analog converter) HAT boards that enable faithful audio reproduction for not much money. IQaudio has been part of the Pi community since 2015 and we've long been a fan of its range of DAC and amplifier boards. Now IQaudio is part of Raspberry Pi and those boards are being added to the range of Pi products.

Reuse old bluetooth speakers and amplifiers with Balena Sound.

Pi-based KVM over IP

If your PC won't boot to an OS, use your Pi as a remote control.

THOSE WHO HAVE needed to access a PC remotely will probably have tried VPN or other applications such as TeamViewer. However, this kind of software only works within the remote computer's OS, which means it can't access the BIOS, reboot, install an operating system or power on the computer. There are several solutions that make it possible for you to remote control a PC independently of its operating system, but using a KVM over IP is one of the most convenient and affordable approaches.

While a store-bought KVM over IP device can be expensive, it's easy to use a Raspberry Pi to create your own. Developer Maxim Devaev designed his own system called Pi-KVM (https://pikvm.org/), which he's planning to sell as a $130 kit. However, if you have the right parts, you can use the software with your Pi, to put it together for far less.

Here, we'll show you how to build your own Raspberry Pi-powered KVM over IP that can output Full HD video, control GPIO ports and USB relays, configure server power using ATX functions and more. You can control the whole setup via a web browser from another device on your local network.

DOWNLOAD THE IMAGE
First, download the SD card image from pikvm.org. There are different versions, depending on which Raspberry Pi you use and whether you use the HDMI-to-CSI bridge or an HDMI-to-USB capture dongle. The file is bz2-compressed, but Etcher can handle this automatically, so follow the guide from earlier and write it to an SD card.

Next, install the HDMI-to-CSI bridge or USB-to-HDMI dongle and prep the OTG USB-C cable.

Connect the CSI ribbon cable from the HDMI-to-CSI bridge to the Raspberry Pi's CSI camera port. Make sure that the blue marking faces the black clamp. If you're using an HDMI-to-USB dongle, connect it to a USB port on your Pi. If you're using a Pi Zero then you'll need a micro-USB-to-USB Type-A hub.

Disable the 5V pin on one of the USB Type-A male connectors from your splitter. The easiest way to do this is to place a small piece of Kapton tape over the right-most pin. If you don't disable it, then it will back feed the power from your wall power to the PC, possibly causing damage to its USB port.

Connect the USB C-to-A cable to the Type-A female connector on the splitter. This will provide power to the Pi. Connect the USB-C cable to the Raspberry Pi 4's USB-C port and connect the unmodified Type-A male to your power supply. Attach the USB Type-A connector and HDMI to the PC that you wish to remote control.

THE PI-KVM SOFTWARE
We're ready to start using the Pi-KVM. On first boot, it will take longer than expected due to the initial process for enlarging the microSD card. Be patient, it will boot eventually. Navigate to the Pi's IP address (look it up in your router's control panel or use Nmap) in a browser on your client computer (the one you're using to control the other PC). You'll be redirected to the login page, which you should log into using `admin` for both username and password.

Click the KVM icon. You should be presented with a screen providing access to the remote PC and a number of other menus. Further options and instructions are available from the Pi-KVM GitHub.

Pi-KVM Prototype V3 without case and KVM four-port switch.

Environment sensors

How does our garden grow? With Raspberry Pis all in a row.

WITH MORE EXTREME WEATHER conditions predicted in the future, the green-fingered amongst us may need to keep a closer eye on their plants. Temperature is one variable that can be displayed and graphed by a Pi with an LCD display, so for this project, we'll take that idea a step further by measuring a range of environmental variables. Thermistors and most other components that measure such factors, are analog in nature, so they aren't immediately suitable for the Pi, which only has digital (GPIO) inputs. However, thanks to HATs like the GrovePi, it's easy to get your Pi talking to analog

ABOVE: Stack the display on the GrovePi daughterboard, plug in sensors and you'll never have dry peonies again.

sensors. We used a high accuracy temperature and humidity sensor, a soil moisture meter, a light sensor, but more are available and getting them working with GrovePi is super easy.

We won't cover setting up the Pi (any model will do), but you'll probably want to get it set up for SSH and Wi-Fi access, especially if you're planning on installing this in your greenhouse. You may have to coerce your Pi into using Python 3 by default (if it doesn't already). This is a matter of

BELOW: Adafruit's Blinka library enables Python to talk CircuitPython, a variety of MicroPython used in gadgets and robots.

```
$ sudo update-alternatives --install /usr/bin/
python python /usr/bin/python2.7 1
$ sudo update-alternatives --install /usr/bin/
python python /usr/bin/python3.7 2
```

Now we can set up the GrovePi board, and Dexter Industries provide a script that will do this in one line:

```
$ curl -kL dexterindustries.com/update_grovepi
| bash
```


VISION ON

You'll need to reboot once everything's done, at which point we can start setting up our display. There are a huge range of TFT, OLED or LCD screens available for the Pi so this depends on your hardware. We

used Adafruit's 1.3-inch Pi Bonnet display and to set that up we had to install Pip and then install its Blinka library, which enables interfacing with their CircuitPython APIs:

```
$ sudo pip3 install --upgrade adafruit-python-
shellwget $ https://raw.githubusercontent.com/
adafruit/Raspberry-Pi-Installer-Scripts/master/
raspi-blinka.pysudo python3 raspi-blinka.py
```

Then reboot and you should be ready to go. Reading from our combined temperature/humidity sensor requires copying and pasting from https://github.com/ControlEverythingCommunity/TH02, since the bit-banging is complex. But our moisture and light sensors can be read with one liners:

```
import grovepi
moisture = grovepi.analogRead(1)
light = grovepi.analogRead(2)
```

Having got data from the sensors, we need to get it on the display. The Adafruit Python libraries enable images to be sent directly to the display, so we can use the Pillow libraries to help us. There is some boilerplate to initialize the display (which we'll omit here, but you'll find it at https://learn.adafruit.com/adafruit-1-3-color-tft-bonnet-for-raspberry-pi/python-stats-example):

```
from digitalio import DigitalInOut, Direction
from PIL import Image, ImageDraw, ImageFont
import adafruit_rgb_display.st7789 as st7789
```

Then you can use regular PIL image and text commands to draw the info on screen

```
height = disp.height
width = disp.width
image = Image.new("RGB", (width, height))
rotation = 90
image.draw = ImageDraw.Draw(image)
```

LOOP THE LOOPS

For this project to be useful, we need the display to continuously update. We make a loop, which clears the screen, updates readings, pauses and repeats

```
fnt = ImageFont.truetype("/usr/share/fonts/
truetype/dejavu/DejaVuSans.ttf", 30)
while True:
    moisture = grovepi.analogRead(1)
    light = grovepi.analogRead(2)
    draw.rectangle((0, 0, width, height), outline=0,
fill=0)
    draw.text((20,120), "Light {}".format(l), font=fnt,
fill="#FFFF00")
    draw.text((20,210), "Moisture {}".format(m),
font=fnt, fill="#00FFFF")
    disp.image(image)
    time.sleep(0.5)
```

You can go further with this project, such as logging and graphing temperatures with Gnuplot. Or connect a watering system via a relay and never have to worry about your peonies drying out again!

Bitcoin Node with RaspiBlitz

Make the Bitcoin network bigger with a Pi and a large storage device.

YOU'RE NOT GOING TO get rich mining Bitcoin on a Raspberry Pi. Even with dedicated ASIC mining hardware, it's hard to compete with industrial operations in countries with subsidized electricity supplies. It's the same for other cryptocurrencies too – see www.tomshardware.com/uk/how-to/mine-cryptocurrency-raspberry-p.

Yet you can still use RaspiBlitz to turn your Pi into a full Bitcoin node and store your bitcoins locally —you're essentially self-sovereign as you haven't entrusted your keys to others and you aren't relying on other nodes to supply information. To quote the RaspiBlitz manual: "Not your node, not your rules".

PUT SOME STORAGE IN PLACE

Running a full node involves regularly syncing the Bitcoin blockchain, which at present is 350GB and growing, so you'll need some large external storage—a 1TB or larger external SSD. Ideally, you want your node to be online all the time, but if you have a fast internet connection, it's easy enough to catch up after some downtime. Bitcoin nodes verify new transactions before sending them to miners, and record new blocks as they're mined, so they can't do anything useful until they're synced.

Using a 3.5-inch LCD touch display to run RaspiBlitz is recommended but you can also run it headless. It's also advised to use a heatsink to avoid damage from overheating. Fetch the SD card image from https://github.com/rootzoll/raspiblitz, write it to an SD card (minimum 16GB), and use it to boot

You'd be hard pushed to find a large-enough SD card to store the blockchain, so get yourself an external

your Pi. If you have the recommended XPT2046 display, this will be used automatically and will display the device's IP address. If not you can use a tool such as Nmap to get this information and then connect from another device with, say

`$ ssh admin@192.168.178.47`

and using the default `raspiblitz` password.

The setup wizard will walk you through setting up passwords, storage, and syncing the blockchain. Once you log in again you'll be presented with the options screen where you can tweak settings and set up as a node on the Lightning Network (see below). Note that for others to connect to your node, you'll need to open port 8333 on your router (and use a dynamic DNS service if you don't have a static IP at home). For Lightning, you'll also need ports 9735, 10009 and 8080 to be open. You can also open a command prompt from the main menu, but don't be tempted to run `apt update` or such. There are instructions for updating RaspiBlitz on the project page and if you don't follow them there will be tears.

THE LIGHTNING SEEDS

Besides running a Bitcoin node RaspiBlitz enables you to run a node on the "layer 2" Lightning network. The Lightning network permits Bitcoin transactions to take place off the main blockchain so they can be confirmed instantly rather than subject to the median six-minute confirmation time that is part of the Bitcoin fabric.

It's possible to generate some income by running a Lightning node, but not necessarily something you should jump into unawares. Since the network lacks the democratic (but lengthy) consensus mechanism of the main network, you could operate a rogue node and attempt to double-spend coins or otherwise cheat the system. However, there are ominous-sounding Watchtower services that detect improperly operating Lightning nodes and penalize them.

Lightning nodes are connected to blockchain wallets, and these must contain sufficient funds to cover any off-chain business. If malfeasance is detected, a so-called Justice Transaction is initiated funds are docked. If your node accidentally goes offline while a channel is open, then you may initially lose funds, but these can be reclaimed (subject to both party's approval) once your node is back online.

You can use Onion (Tor) routing to keep the IP address of your Lightning Node a secret.

MotionEyeOS

Build your own motion-triggered home security camera.

HOME SECURITY CAMERA SYSTEMS have exploded in popularity and decreased in price over the past few years. Yet there are some drawbacks: vendors often charge a monthly fee to store your data; and you might not want to share video and photos from inside your home with a third party.

MotionEyeOS is a distro that enables you to turn a Raspberry Pi with a camera into a home video-monitoring system, where the photos and videos can either stay on your device (and home network) or, if you choose, be uploaded automatically to a cloud-storage service such as Google Drive or Dropbox.

Here, we'll show you how to set up a Raspberry Pi security camera with MotionEyeOS. This software works with almost any Raspberry Pi (connected to the internet) and most webcams or Pi cameras too. There's no fancy coding in this project; it just works.

Download the latest version of MotionEyeOS corresponding to the model of Pi you're using from https://github.com/ccrisan/motioneyeos/releases and write it to the SD card. When the process completes, physically remove and then reinsert your microSD card. We do this because the software automatically ejects the microSD card when the process completes, but we need to add one file before the next step.

Create a file named wpa_supplicant.conf with the this text, replacing "YOUR_NETWORK_NAME" and "YOUR_NETWORK_PASSWORD" with your details:

```
country=gb
update_config=1
ctrl_interface=/var/run/wpa_supplicant
network={
 scan_ssid=1
 ssid="YOUR_NETWORK_NAME"
 psk="YOUR_NETWORK_PASSWORD"
}
```

This Raspberry Pi security camera can be used to monitor children or pets, or to watch out for burglars.

Save the file, eject the SD card and insert it into your Pi. Connect your camera, monitor and power supply to your Raspberry Pi and power up. If you have a monitor connected your Pi's

Don't rely too much on SD cards for storage. Cloud saves will make files easier to access too.

internal IP address will be displayed on the Pi screen. Alternatively, you can find this out from Nmap or your router's configuration page.

Enter your internal IP address into a browser window and MotionEye should start streaming.

CONFIGURING MOTIONEYE

Click the Profile icon within your browser menu to pull up the Login screen. Log in using the default credentials (username is admin, the password field should be blank). Select your Time Zone from the menu in Time Zone. Click Apply. MotionEye will reboot which will take a few minutes. This is important as each photo and video is timestamped.

The Frame Change Threshold setting determines the proportion of pixels that change before recording starts. Set your percentage low enough to pick up the movement you're tracking, but high enough to avoid recording a passing cloud. Start with the default 4 percent Frame Change Threshold and then move up until you reach your optimal setting. Click the down arrow to the right of Still Images to reveal extra settings. Do the same for Movies. Set Capture Mode and Recording Mode to Motion Triggered and choose a period to preserve both. A week is reasonable if you're working with an 8GB card. Click Apply to save changes.

Set your Camera Name, Video Resolution, and other options in the Video Device section. Click Apply to save changes. On the live feed view, you'll see buttons for viewing saved images and video.

In the current setup, this media won't leave your local network. But if you're okay having your media in the cloud, it's easy to set MotionEyeOS to upload to Google Drive or DropBox. If you want to enable remote access to your Pi security cam, you'll need to configure port forwarding on your router. If you do this, there are official iOS and Android apps so you can check your home when you're out. ⏻

CENTERFOLD

1 INTERNAL SPECS

Each mesh router comes equipped with a 1.5 GHz tri-core processor, 256MB of Flash storage, and 512MB of RAM. On top of that, they both pack one RJ45 WAN port and three RJ45 Ethernet ports too.

Asus ZenWiFi XD6

DESPITE OUR RAPID progress as a species over the last 200 years, the global footprint for radio noise has dwindled in the past two decades. Radio waves are slowly becoming quieter, the very things that ushered us into the realm of unhindered communication all those years ago are now becoming more refined and more sophisticated.

No longer are we churning out huge gobbets of information and throwing it deep into the dark abyss of space. No, the inevitable march of technology has brought ever more advanced forms of communication that have defined this era. As our computational devices grow more potent and we produce deeper and richer content, the need for more impressive forms of compression and greater bandwidth has become paramount. From 2.4 GHz to 5 GHz, and now Wi-Fi 6 and 6E, all have improved upon the past generation, providing us with greater capacity, and all of them have become quieter as they do it.

Interference then becomes king, waves that could pass data through walls, although still here, are now choked and humbled, but while their successor's bandwidths are far more powerful, they are hindered by such things as raindrops, or even the hum of a microwave oven.

While our homes haven't changed quite as rapidly, the need to boost the future of wireless has become a paramount part of the tech enthusiasts' place of rest. Enter the Asus ZenWiFi XD6. These twin routers can bathe up to 5,400 square feet in the brilliance of the modern era, inside and outside your home. Let's take a look. **–ZAK STOREY**

2 DESIGN

From a design standpoint, the clean, crisp, white style with silver and gold elements gives the ZenWiFi XD6 a subtle appearance that means it should blend in with its surroundings. This isn't a ROG flagship product, laden with antennas and lights, and it benefits from that.

3 WI-FI 6

At the heart of the ZenWiFi setup is support for Wi-Fi 6. This new and improved standard is up to 2.2x faster than the previous Wi-Fi 5 routers and can pack a peak transfer rate of 5,400Mbps. That's a huge amount of data, but it does require you to be within relatively close proximity, making these mesh routers vital to its application around your home.

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

SAM LEWIS
STAFF WRITER

HOW LONG TO BEAT

If you love video games and planning ahead then this might just make your day. How Long To Beat (https://howlongtobeat.com) is a website that gives you average completion times for a vast library of games on both the PC and game consoles past and present. It's a great place to head to before diving straight into a game, so you know whether you can fit it into your schedule or need to book a whole month off work to complete it.

OS EXCITEMENT

Let's set things straight here, regardless of the compatibility issues that have been brought up by the Windows 11 announcement, I have been loving testing out the new operating system. I signed up to the Windows Insider Program, which makes you feel like you're on some undercover SWAT operation, when, in reality, Microsoft is using you as a guinea pig. Swings and roundabouts, right? I chose the Dev channel and downloaded the latest version of the pre-release Windows 11 build, which in layman's terms means I will get all the bugs. Thankfully, I haven't stumbled across many, apart from the Wi-Fi icon in the taskbar not being centered correctly. However, this seems to have been patched, and hasn't happened after a restart. I'm now questioning whether it was a bug, the patch, or if I need new glasses.

Going back to the topic of centralization, I personally am a big fan of the dock being in the middle. Coming from the Apple ecosystem and having used Macs for previous jobs, it's just something that feels natural to me. Those who don't like it can always align it back to the left. The design brings a whole new edge to the operating system and, seeing as I'm enjoying this early build so much, I'm excited for the official release. Let's just hope they make it more compatible.

MAKE – USE – CREATE

58
How to create 3D typography in Adobe Illustrator

64
Our step-by-step guide to installing Windows 11 fresh

66
Building high-end ITX art with a G.Skill Z5i chassis

© MICROSOFT

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

IFIXIT presents:

AUTOPSY

PS5 DualSense Controller

Nothing says teardown like a snazzy X-ray eh?

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.

The DualSense controller turns out to be among the PS5's biggest upgrades.

Alien tech, or the latest brilliance in analog controls?

BACKGROUND

While we were waiting for our PS5 to arrive, we got busy on the DualSense controller. There's enough new stuff packed in to make this a meal on its own:

MAJOR TECH SPECS

• Space-age adaptive triggers

• Top-of-the-line haptic system

• SIE CXD9006GG—likely a custom Sony chip doing all the work

• Dialog DA9087 PMIC

• Realtek ALC5524 audio codec

• Nuvoton NAU8225 3.0 W Class-D Audio Amplifier

• 5.7 Wh Li-ion Battery

KEY FINDINGS

• The new design is so sleek that there are zero visible screws. We aren't fazed though—where there's a seam, there's a way! The black trim cover around the joysticks unclips with ease, revealing two screws near the tips of the handles. Two more screws are hidden under the L1 and R1 buttons, which come up with a little clip-prying (and potential flying).

• Lift off! Four Phillips screws and some clips are the only things standing between customers and an open DualSense controller. Just like the DualShock 4 and DualShock 3, the DualSense battery has a tough plastic shield and isn't glued in place. For battery-swapping serviceability, this beats almost any modern smartphone. This grey monolith clocks in at 5.7 Wh—a sizable increase from the DualShock 4's 3.7 Wh pack. There's a reason for that extra juice: something has to power all this tech. The new DualSense cell is more in line with the 5 Wh pack in Nintendo's Switch Pro Controller, which is also easy to replace.

• Next up, out come the guts: button sensors, motherboard, haptic drivers, and adaptive triggers, all mounted to a black midframe. From this perspective, it's clear how much faith Sony has in the DualSense's upgrades. Half the internal volume is dedicated to the triggers and haptics! Both adaptive trigger assemblies connect with cables, but the haptic drivers, USB-C port, and joysticks are all held down by soldered connections.

• Let's dig into this classified alien tech... er, variable resistance trigger [**Image A**]. The trigger works normally without any of the mechanics active, making contact with a button on the blue ribbon cable. But game developers can choose to program the controller to precisely adjust how difficult the trigger is to pull. The silver motor spins the white worm gear, which drives the black lever arm up and provides resistance to the trigger's lever action, adding another level of feedback to the controller.

• With all the easily replaceable parts out, we're down to the soldered-on bits. First off are the two Foster-branded voice coil actuators driving the haptics system. Next: the two joysticks. These look identical to the Alps-brand sticks from the DualShock 4. Compared to these, a Joy-Con joystick replacement is a breeze. PlayStation joysticks are less prone to drift than the notoriously drifty Joy-Cons, but soldering them on is a big miss. Finally, the USB-C port: another high-wear component that we'd prefer to see easily replaced.

• The DualSense controller turned out to be among the PS5's biggest upgrades. As a final treat, here's an X-ray look at all the DualSense we just dismantled [**Image B**] thanks again to our friends at Creative Electron. We also owe a huge thanks to iFixit contributor Chunglin Chin, who did the detective work needed to ID all those chips in the motherboard steps. Great job!

• Repairability Score: 7 out of 10 (10 is easiest to repair). ⏻

3D Typography in Illustrator

IF YOU HAVE FOLLOWED ALONG with the first Illustrator tutorial in the previous issue, you should have picked up a few of the basics within the application and, hopefully, can work your way around it. Like we said before, if you are familiar with the Adobe ecosystem and have used other pieces of software other than Illustrator, don't sweat too much as it should be fairly easy to pick up. These applications work hand in hand with one another and are pretty similar.

One of the main uses for Illustrator is for graphic design, and one huge sector of that department is typography. I hope we haven't lost you already, trust us, you can create some nice artwork using type, especially with the handy help of Illustrator. One particularly useful exercise is wrapping text around shapes. Using this skill, you can create tons of attention-grabbing text for your design, artwork, advertisement, or whatever you intend to do with it. It's tips such as these that can set your style apart from the typical flat text that you see everywhere. Hopefully, with this skill in the bag after this tutorial, you can piece it together with another project to give it more fizz. You can't go wrong with that, right? –SAM LEWIS
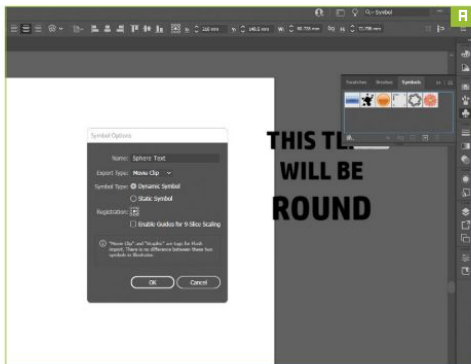
1 SPHERICAL TEXT

To get started with some text wizardry, we need to create a new artboard by opening and creating a new file of your size and color mode. Remember, CMYK is for printing and RGB is for digital files, so pick accordingly. Also, remember to name your file too as it helps to stay organized working on projects. Once all the necessary mundane setup is complete, we need to add some flat text to our work. You can either press the 'T' key to open up the type tool or head over to find the type tool icon on the left-hand side. Then you can either click on the artboard or draw out a text box and type what you want. We thought we would be hilariously ironic with this one. Please accept our apology for the awful sarcasm and proceed with the tutorial...

›› For this part of this tutorial, we will be focusing on creating some 3D spherical text. Once you have typed out your text, you can also play around with it first using the preferences towards the top of the application [**Image A**]. These include, color, font, font style, text size, alignment, and so on. Then, once you are happy with it, open the symbol preference tab from the right-hand tool section on the app and simply drag the text box into it. The symbol tool is the same icon as the clubs suit from a pack of playing cards, which makes it a little easier to find.

2 MAKING 3D TEXT

After creating a symbol, we can remove the text from the artwork by deleting it. No, we aren't finished yet, if only it was that simple! In fact, this is where the wizardry begins. We now need to draw the shape from which the text will take its form. Head over to the left-hand side and click on the shape tool, select the ellipse tool and draw a circle on the artboard. We actually don't need the full circle, just a semi-circle so then once the circle is drawn, use the scissor tool to click on two anchor points and split the shape in two. Delete one side and you should be left with a semi-circle like so.

›› Now we are going to delve deeper and make this 3D. Head over to the effect tab at the top, 3D, then the revolve option, this will open up a new window. Once in this window, click the map art button to open up another window. Before we get too confused, let's just take a second to focus on this new pop-up window. Yes, your screen does look like it's straight out of a TIE fighter cockpit, but it's not as complicated as it seems. In this new window, check the invisible geometry box to make sure the sphere is clear [**Image B**]. Then at the top, we can see a symbol drop-down menu. Remember earlier when we created a symbol out of our text? Well it's all starting to piece together now, isn't it?
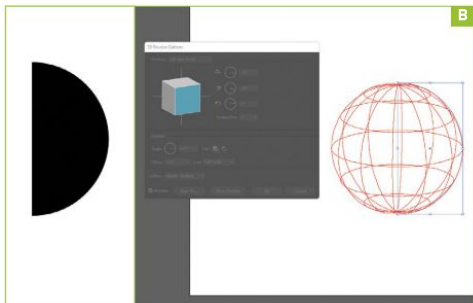
WHY TYPOGRAPHY MATTERS

When it comes to design, there are a few things that need to stand out or work together in order for it to be effective.

First, you have the actual main artwork, image, or design. This is the focal point, for example, if you are creating a poster. Then you also have the color scheme for the design. Don't go overboard with this as a lot of colors together can clash and take away from the main design, unless that is your intended style. Last but not least, we have typography. This also has to stand out and not get lost in the design. However, you don't want to necessarily allow it to swamp the whole piece.

This tutorial shows you how to can use type as an image, which is a great way to combine the two and create a great hierarchy. This is also something that is extremely important in design. Hierarchy creates a path for the consumer's eyes to follow. With this in mind, figure out what you want your audience to see first. This is usually the most important part of the design, for instance, an image or the name of something. In typography, you want this to stand out the most over other areas of text. In most cases, this will be the title.

For an event poster, you have a title, then a strapline, maybe some features of the event, and at the bottom, in smaller print, usually all the key details. There are certain styles in typography that help you create a hierarchy. You don't necessarily need to change the font to create a hierarchy, and in most cases, you should stick to a maximum of two fonts in your design.

When you have chosen a font, you can change its style to create differences, for example bold, italic, underlined, heavy, or light. You can also change the font size, of course, and the color can be changed too. There are spacing changes you can make to increase or decrease the room in between the lettering (kerning) and alter the spacing between characters in a block of text (tracking).

There are plenty of ways you can do to draw attention to different parts of your text other than just the size of the text. Have a go at creating your own text hierarchy and play around with all the typographic settings to create exactly what you want. Use this tutorial to create imagery out of your text and produce a striking graphic that combines both design and typography skills.

3 ADDING SYMBOLS TO THE SHAPE

As you may have guessed, we need to click on the symbol drop-down menu and add our text symbol straight into this shape. This may look a little odd as soon as you place it in, but this is where we can do some tweaking. In the curved line graph chart, we can scale the symbol (text) and move it around to the desired position. Quick tip, in most pop-up windows in the Adobe suite, if it has a preview check box in the bottom corner click this to see the preferences happen in real-time. Make sure this is checked to see how your text will look with a 3D spherical effect. [**Image C**]

›› With a little playing around, you can get this exactly how you like and then hit the ok button when you are happy with it. Then in the 3D resolve options window, we
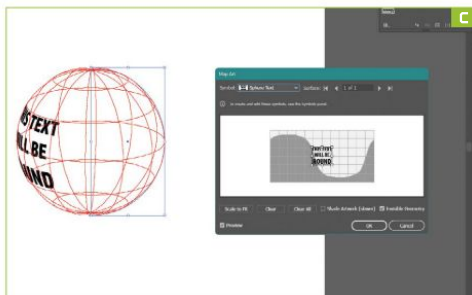
Join the conversation!

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

forums.pcgamer.com

PC GAMER
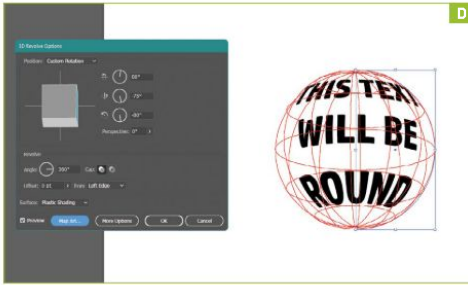
can play around with the XY&Z values to pivot the sphere further to get the correct angle you want [**Image D**].

›› For good measure, we have added a colored circle in behind the sphere to give it a ball-like effect and here is the finished project [**Image E**]. Using this technique could come in handy for creating product mockups where you have rounded text on a spherical object, such as a ball or a globe.

4 CREATING CUBED TEXT

Right, so we have done spherical text, now let's add a few right angles and try to create cubed text, sound good? Presuming you have said yes, we will continue with this part of the tutorial. For the cube text effect, we need to start by creating the text again. For this instance, we are going to use the same text over all three faces of the cube that we can see. To make things more interesting, you can create three different pieces of text with various typographic changes to create a hierarchy for your text. Once you have your type ready, drag and drop it into the symbol tab once again. Then we need to create a square which we will be turning into a cube. Go to the shape tool, find the square shape, and draw it out.

5 MAKING A 3D CUBE

To make this 3D, head to the effects bar at the top, then go to 3D and choose the extrude and bevel option this time, instead of the revolve option. This will create a cuboid instead of a perfect cube but depending on your text, it will still work. If your text is square-shaped like ours, it makes sense to create a cube so we need to play around with the extrude depth values in this menu to create this. This will give you more or less depth depending on whether you increase or decrease the value. We chose a depth of 260pt, which gave us a decent-looking cube.

›› Once happy with your cube/cuboid, we need to open up the map art window again by clicking the button at the bottom of the 3D extrude and bevel options window.

6 INSERT SYMBOL HERE

Inside the map art window, we can insert our symbol once again. First, though you need to determine which surfaces you will need, for the purposes of this tutorial, you want to be using all three facing surfaces [**Image F**]. In our example that was surfaces 1, 5, and 6. Once on the surface, add your symbol (text) by going into the drop-down menu and selecting it. The little graphic in the center of the pop-up window will show you how this will appear.

›› To make things look neater, we used the scale to fit option to fill the face of the shape but then scaled the symbol down a little to leave a clear gap around the edge of the text. On the top and side surfaces, we applied the same method but had to rotate the symbol for it to read correctly. Then tick the checkbox that says invisible geometry to make the shape transparent and once you are happy with the layout and design, click ok.

7 BUILD YOUR OWN DESIGNS

You should hopefully now have a 3D cube/cuboid shape that somewhat resembles what we have designed. We say somewhat, as it is good to practice loosely on our steps to create unique pieces of work. These tips should be taken as building blocks for your own designs and to be used in conjunction with other techniques to create something amazing. Hopefully, you can take something from this tutorial and keep tweaking it exactly how you like [**Image G**].

›› Illustrator is all about taking ideas from your head, or even sketches on your paper and transforming them into brilliant designs and artwork. It takes time and effort, but once you get into the swing of things and learn the ins and outs of the software, it will become second nature. Then you can start combining lots of different ideas to help boost your portfolio of work.⏻

Machine of the Month:
Original Macintosh (1984)

THE ORIGINAL MACINTOSH (or "Mac" for short) is an iconic machine that brought home desktops to the masses and established a line of computers that is still going today. Its compact form, cool styling, and elegant graphical interface immediately won the hearts of creative people worldwide, and would influence many desktops that followed.

However, it's now very difficult to find original Macs in good working order, and it's even harder finding working Macintosh floppy disks.

This leaves emulation as the only way most people can now experience these machines, which can be tricky in itself. So, after a brief history tour, let us guide you through the process of emulating the Macintosh. **–JOHN KNIGHT**

1 DEVELOPMENT AND RELEASE

Launching in January 1984, the Macintosh was the first successful home computer to use a graphical desktop, and also to use the 16-bit Motorola 68000 CPU. This followed 1983's Apple Lisa desktop workstation, which also had these features, but was too expensive and sold poorly.

›› The Macintosh was originally the brainchild of Jeff Raskin, who envisaged a much cheaper machine, but was taken over by Steve Jobs after being forced out of the Lisa team.

›› The Mac's operating system, System 1, was a cut-down version of the Lisa's desktop. Both took design elements from Xerox workstations (the first machines to use a graphical desktop) in exchange for Xerox owning Apple shares.

›› The Macintosh was a near opposite to Apple's main seller, Steve Wozniak's Apple II. Where the Apple II valued openness and functionality, the Macintosh was extremely locked down and valued form, packing the entire system into a small carry-box with an integrated 9-inch monitor. Surprisingly, the Macintosh was black and white, while the older Apple II was color.

›› The Mac was an exercise in minimalism: the mouse had only a single button, and there were no arrow keys or function keys on the original keyboard. Apple wanted to encourage customers to use the mouse, and for developers to design applications around the desktop.

›› Standard units shipped with a single 400K 3.5-inch floppy drive and no internal storage. This resulted in a lot of disk swapping, but external hard drives and floppy drives were available for the impatient.

›› Despite its multi-window interface, the Macintosh was originally a single-tasking machine, though this changed with the introduction of Switcher in 1985, and 1987's System 5 OS came with multitasking out of the box.

2 SOFTWARE

Macintosh gaming had a large emphasis on the mouse. Many of these games are quite shallow, but there are some standout titles, like *Shufflepuck Café* (1988) and *Pararena* (1990), and strategy games benefited greatly from this new mouse input.

›› Trivial mouse games aside, early Macintosh gaming is an interesting experience. Eighties gaming was obsessed with color, but using color requires more RAM. Original Macintosh games were restricted to black and white, which allowed higher resolution gaming that was often much sharper and more detailed than other platforms of the time.

›› *Dark Castle* (1986) and *Prince of Persia* (1992) are two examples that feel very different, and arguably better, than rival systems' colorized equivalents.

›› But it's really in desktop applications where the Macintosh shines. Macintosh software is simple and elegant. Try something like MacWrite and it is much simpler, quieter, and more intuitive than today's bloated word processors. The original Adobe Photoshop is a joyful thing, and MacPaint makes a mockery of Microsoft Paint. If only it were color.

3 LEGACY

Year by year, the original Mac underwent many evolutionary changes.

›› The Macintosh 512K (launched in September 1984) quadrupled the system RAM. The Macintosh Plus (January 1986) and SE (March 1987) models added 1MB-4MB of RAM, and space for an internal hard drive or second floppy drive. 1989's SE/30 came with a high-density floppy and RAM expansion up to 128MB. Finally, the stripped-back Macintosh Classic sold from 1990 to 1992 as a budget machine.

›› It's difficult to find sales numbers for the original Macintosh, but though initial sales were stronger than expected, those numbers soon decreased and it took until 1987 to sell a million units.

›› **Although late to the party, 1992's *Prince of Persia* showed how much further old machines could be pushed if only more developers were willing to try monochrome gaming.**

SPECIFICATIONS

ORIGINAL MACINTOSH (1984)

CPU:	Motorola 68000 @ 7.83 Mhz
RAM:	128 – 512 KB
Graphics:	Black-and-white @ 320x256
OS:	System 1 - System 7
Storage:	3.5" floppy, 400 KB - 1.44 MB, plus hard disk upgrades
Launch price:	$2,495
Released:	January 1984
Production:	1984 - 1992

>> **In an early example of stripping useful features, Apple removed arrow and function keys from the Mac's keyboard. But at least it had a headphone jack.**

>> The original Macintosh enjoyed a production run from 1984 to 1992. It was eventually overtaken by its successor, 1987's Macintosh II series, featuring color graphics and a 32-bit CPU.

4 EMULATION

Initially, we found Macintosh emulation quite difficult, but we've streamlined the process the best we can. The two main emulators are Mini vMac and Basilisk II. Advanced users should try Basilisk II, but we found Mini vMac easier to get started and more stable, so we'll use that here.

>> Before emulating the Mac you need three things: a system ROM, an image of Macintosh OS, and a copy of StuffIt Expander – everything in Mac world uses StuffIt compression.

>> Legally, we can't recommend system ROMs or OS images unless you own the original products. But if you are in the clear, archive.org has a collection of Macintosh system ROMs. The MacPlus v2 and v3 ROMs are the only ones that work well with Mini vMac.

>> As for the OS, the Mini vMac website has links for System Software 6 boot disks, though annoyingly they are double compressed, so you will have to extract the file twice until you're left with a file called "System Startup". Otherwise, if you want System 7, macintoshgarden.org has a software entry entitled "Mac OS 7.5.3 (for emulators)".

>> System 6 is simpler and boots a lot faster. System 7 is a lot more cluttered and slow to boot, but also more compatible. If anything doesn't work in System 6, try System 7.

5 GETTING STARTED

Once you have all the prerequisites, head to the Mini vMac website and open the Download page.

>> For anyone who is interested specifically in the first models, there is a 128K emulator variant, but for everyone else, click the link 'Download Standard Variations'.

>> Compressed files are provided for Windows, Linux, OS X, and other platforms. The Windows ZIP file contains a single executable file called Mini vMac.exe. The Linux package contains a single executable file that you may be able to run just by

clicking on it. Otherwise, you can open a terminal and run it by entering:

```
$ ./Mini\ vMac
```

>> Once started, Mini vMac will give you an error message, warning you that it can't find a ROM image. Drag your chosen system ROM onto the Mini vMac window, and if your ROM is recognized there will be a loud beep and a floppy disk icon. This is a normal Macintosh screen – it wants a system boot disk. Click and drag your OS image onto the window and, all being well, your OS should boot.

>> Once your OS has loaded, you can explore disk images simply by dragging them onto the Mini vMac window, and if the image mounts successfully, a disk icon will appear on the desktop. If Mac OS says the disk image isn't readable, click Eject. Do not click Initialize – you may wipe out the image file!

>> Mini vMac has strange default settings and weird controls. The emulator is controlled by holding down Ctrl then pressing another key.

>> Firstly, Mini vMac runs at 8x speed by default. To change the speed, press Ctrl + S, which will bring up a speed menu, while keeping Ctrl held down, then press Z. The emulator runs in a small window by default, press Ctrl + M to make the window bigger. To make the window fullscreen, press Ctrl + F.

>> Be aware that you will make changes to your OS image, and it doesn't like being shut down badly. To shut it down properly, open the Special menu from the top menu bar and choose Shut Down, after which you can safely close the window. ⏻
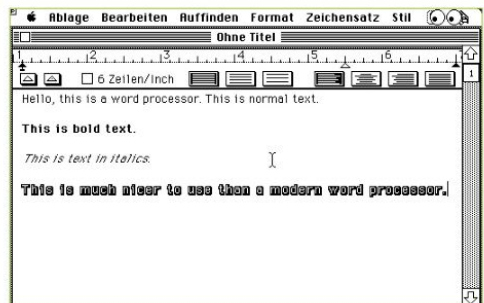
>> **If you've ever had a migraine just looking at Microsoft Word or LibreOffice, MacWrite offers you clean and simple serenity instead.**

How To Install Windows 11 Fresh

YOU'LL NEED THIS

8GB USB 3.0 FLASH DRIVE
Internet Connection.
Another PC or Laptop

THE PRODIGAL SON HAS ARRIVED, Windows 11 is here in all of its insider glory, complete with quirky new taskbar layout, updated start menu, and enough options-baked-into-options to make even the most enthusiastic PC nerd happy (seriously, this particular Editor-in-Chief is overjoyed with it all). That said, in the world of Windows you often get two camps of people—those who like to perpetually update their machines, never performing a clean reinstall no matter the bugs or issues, and those folk opposite to that. The ones who will happily flatten everything at the drop of a hat if it means a cleaner system.

And let's face it, with the myriad of problems that occurred moving from Windows 7 to Windows 8 and 8.1, and then from Windows 8 to Windows 10, the idea of a fresh install is certainly appealing. So, as we're still waiting on Microsoft to release an official ISO or Media Tool Kit for Windows 11, we're going to need to jump through some perfectly safe hoops to get ourselves a squeaky clean install of the latest OS for you to play around in. –**ZAK STOREY**

1 COMPATIBILITY CHECK
Now, technically, you can just install Windows 11 from ISO right now without too much of a worry—you don't need to enable Secure Boot, nor do you need to enable TPM 2.0. However, this may change in the coming months after various updates kick in, so the first thing you'll want to do is identify whether your system is compatible or not.

›› We've covered this in-depth across multiple features and columns at this point, but basically, you're going to need either an AMD Ryzen 2000 series chip or higher, or an Intel 8000 chip or higher to get this sweet, sweet Windows 11 goodness on your rig.

›› If you can tick both of those boxes above, the first thing you'll want to do on the system upon which you're going to install Windows 11 is head straight into the BIOS. Once there, you need to be looking for two settings, typically found in the advanced area and boot areas of your UEFI BIOS.
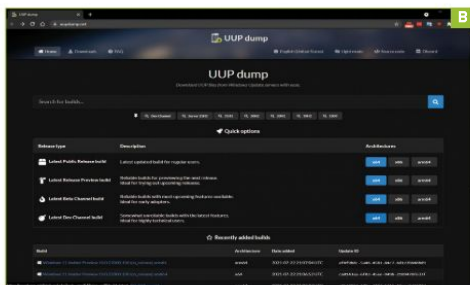
GET YOUR DRIVERS

Before you pull the trigger and start installing Windows 11, we'd recommend grabbing your drivers. There's nothing worse than installing a fresh new OS only to realize your ethernet ports no longer work, or your GPU driver has defaulted to 800x600 on your 34-inch superwide screen.

Head to your motherboard's product page, and grab a few basics, including the Ethernet and Wi-Fi drivers, sound drivers, and chipset (the rest you can get later if you need them). One last thing to note is the graphics driver: currently, only Nvidia has officially released drivers with support for Windows 11. That said, the majority of the operating system is based on Windows 10, so program and driver incompatibility should be fairly minimal.

We've installed everything we'd usually do for a test build with this tutorial, and there haven't been any compatibility issues, apart from a few UI bugs, something we expect with an insider dev build.

›› Their location will vary depending on the motherboard and processor, but first up you're looking for "fTPM configuration" [**Image A**]. Once in there, make sure the option for "Firmware TPM" is selected. Then head into the "Boot" section of your BIOS, and look for "Secure Boot". Head into that section, make sure it's enabled, and that the OS Type is selected as "Windows UEFI mode". Then, it's simply a case of hitting F10, "Save All and Exit", and you're good to go.

2 FINDING THE ISO
As we mentioned earlier, Microsoft hasn't released an ISO or Media Tool Kit for Windows 11 just yet, so we're going to need to head to a third-party website to download one. In our case, we're going to UUP dump https://uupdump.net. It's all legit, and cleared, and is a community-managed repository for Windows updates in general.
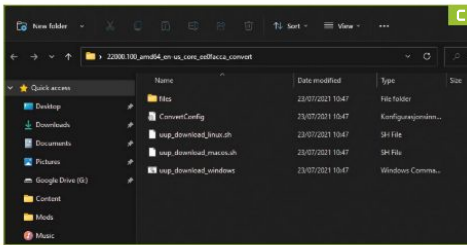
>> Once on the site, you'll notice there's a list of available builds and updates available [**Image B**]. The one we're looking for is "Latest Dev Channel build". Look to the right of that and you'll have three options. One for x64, one for x86, and one for arm64 (yep they even have the ARM variant). Click the x64 link and it will take you to the latest update. It should say something such as "Cumulative Update for Windows 11 (10.0.22000.100)" or similar. Click the blue hyperlink and it takes you to the next page. Make sure you choose the right language for you (in our case that's "English (United States)"), and then hit "Next".

>> At this point, you'll be able to select which version you want to test out. You can choose between Windows Home, Home N, Pro, Pro N, and Windows Team, and although you can opt to download all of them, we recommend that you only pick one. We selected Windows Pro, as that's the license we have tied to our Microsoft Account.

3 CREATE THE ISO

On the next page, make sure that the download method is selected as "Download and convert to ISO", and that under "Conversion options" the checkbox for "Include updates (Windows converter only)" is selected as well. Then, hit the big blue "Create download package" button. You'll get a small zipped folder labeled with the version number, plus the architecture you're using (amd64 in our case). Extract that to your desktop.

>> Open the folder and inside you should find four separate files and a folder called "files". There should also be three additional "uup_download_X" files in there, one with a linux.sh suffix, one with a macos.sh suffix, and lastly one with a windows suffix. Double click the "uup_download_windows" suffix, Windows might tell you it's a dangerous file to run, hit "more info" then "Run anyway" [**Image C**].

>> A command prompt will open up, asking for admin permission, allow it, and it starts downloading all of the required ISO files for you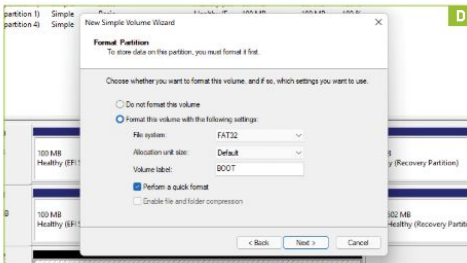 directly from Microsoft's servers. Leave it to complete, it can take a while depending on your internet speeds, and at certain moments may look like it's hanging (it's not, just wait it out). Once that's done, it'll place an ISO file in the folder.

© MICROSOFT

4 CREATING A BOOTABLE USB STICK

Now we need to make a bootable USB stick. This is a workaround, as we can't mount the ISO directly to the USB drive. Most ISO mounting software automatically formats the drive to NTFS, while we need it to be FAT32 to boot from (Secure Boot doesn't allow booting from an NTFS drive). The problem is, some of the ISO files are over 4GB, meaning they can't be stored on that FAT32 partition. To get around this, we're going to have to format our USB stick in a particular way.

>> Plug your USB drive into your machine, then go to the start menu and type "partitions", and open up "Create and format hard disk partitions". Find your USB stick and delete any partitions that are currently on the stick (this will delete all the data on it, so back it up first).

>> Once complete, right-click the unallocated space and select "New simple volume". This is going to be our FAT32 boot partition, so you'll want to make it 1000MB. Set the file system to FAT32, name it something memorable (in our case BOOT), [**Image D**] then perform a quick format. Do the same with the remaining space, but make sure it's NTFS and again label it something different "DATA" as an example.

5 MOUNTING THE ISO

Go back to your Windows 11 ISO download folder, and locate the ISO file it should be labeled as "22000.100.210719-2150.CO_RELEASE_SVC_PROD2_CLIENTCORE_OEMRET_X64FRE_EN-US" or similar. Right-click the file and select "Mount", this will create a virtual CD drive and mount the ISO file to it, giving you access to everything inside. Select and copy everything there, except the "sources" folder, and paste it into the "BOOT" FAT32 part of your USB stick we created earlier.

>> After that's done, in the same FAT32 partition of your USB drive, create a folder called "sources". Now go back to your mounted ISO folder, into the "sources" folder there, and copy the boot.wim file, and place it in your new "sources" folder on your BOOT/FAT32 partition.

>> Finally, go back to your mounted ISO, copy everything, as well as everything you copied before, and paste it all in your DATA/NTFS partition on your USB stick. And there you go, one complete bootable USB stick with a Windows 11 ISO on board ready for a fresh install.

>> Now it's simply a case of plugging the USB stick into your machine, booting off it, following the usual install procedure and you'll have a slick Windows 11 desktop.⏻

BUILD IT

ZAK STOREY, EDITOR-IN-CHIEF

High End ITX Art

G.Skills to Pay The G.Bills

LENGTH OF TIME: **2-3 HOURS**   LEVEL OF DIFFICULTY: **MEDIUM**

THE CONCEPT

IT'S FAIR TO SAY that when we first set eyes upon the G.Skill Z5i in its glorious review photoshoot, we were impressed. Even the photographer turned around to us and said: "I've shot a lot of big old PC boxes for you guys, but this thing is nice". And we have to agree with that sentiment.

On first appearance, the Z5i is an exceptionally beautiful chassis, complete with an aluminum frame, curved tempered glass panels and baked in RGB lighting, including a neat little underglow that makes us think back to the glory days (if you can call them that) right at the start of the *Fast and Furious* franchise.

And it got us thinking about what exactly to build inside this thing? We wanted something a little special, a little different, and a bit more high-end. So we decided to grab a hold of our AMD Radeon RX 6800 XT, gather together a few choice parts from manufacturers across the globe, and set to work, with the aim of building a fantastic ITX work of art. Or at least, that was the plan.

Of course, getting parts right now is a challenge and something that doesn't seem to be changing any time soon. So, we appreciate it's not exactly realistic to build this rig, not least when a $700 GPU is still being listed at over double its RRP price, and for a stock reference card at that. Frustrating? Very.

THE LIGHT FANTASTIC

AS WE MENTIONED, this build was and is all about that chassis. The PC and hardware itself inside, are relatively insignificant in the grand scheme of things right now, so consider this more an exercise on peripherals and case-building, or an in-depth review, rather than a "we built this for X, look how it performs" kind of scenario. If you're interested in flashy-looking ITX cases that can sit proudly on your desktop, or how G.Skill's first attempt at case design went, it's definitely the build-it for you.

That aside, the hardware contained inside, at reasonable prices, would make a fairly chunky gaming PC, or a mid-range entry point for anyone looking to get a foothold into video-editing for sure. The main thing that decided our picks for this issue though? Lighting. The Z5i, being the showcase chassis it is demands some top-quality hardware, and to that end, we went into this with several components in mind.

For our RAM, Dominator Platinum RGBs from Corsair were the logical choice, the cappelix LEDs situated in those things are insanely bright, and once configured properly, can light up a room. Our CPU cooler is the digital display equipped NZXT Kraken Z53 240mm AIO (this case supports up to a 280mm AIO, which we would recommend you go for). And lastly, our GPU is the mildly flashy Radeon RX 6800 XT from AMD.

Originally, we intended to go for one of Zotac's GeForce RTX 3080 Ti's complete with over-the-top ridiculous LED lighting (as Zotac is one of the best at that), but sadly the card arrived a day after our shoot. Don't worry though, the availability and pricing would have still been nonsensical for that card. Oh, boy... this industry is frustrating right now.

INGREDIENTS

PART		STREET PRICE
Case	G.Skill Z5i ITX Tower	$200
Motherboard	MSI MPG B560I Gaming Edge WiFi	$170
CPU	Intel Core i5-11600K	$270
GPU	AMD Radeon RX 6800 XT	$1,300
Memory	32GB (2x16GB) Corsair Dominator Platinum RGB @ 3200 MT/s	$255
PSU	750W Corsair SF 750 80+ Platinum	$182
Primary Storage	1TB Corsair MP600 Force PCIe 4.0 M.2	$160
CPU Cooler	NZXT Kraken Z53 240mm AIO	$230
OS	Windows 10 Home 64-bit OEM	$32
Total		$2,799

1 PRODIGAL CHASSIS

AND HERE WE HAVE the G.Skill Z5i, unboxed and ready for building in. We'll list a few key specs here; it has curved tempered glass panels on both sides that swing open, with some hefty magnets behind them. It supports up to a 280mm radiator (and two either 120 or 140mm fans), and can theoretically support three-slot graphics cards up to 12.99-inches or 9.05-inches with a 3.5-inch HDD installed. It can also support CPU cooling up to 2.75-inches, has multiple support for drives (but we'd recommend you only use M.2 drives), and fits SFX power supplies only, although again the smaller the PSU you can get (as in physical size), the better.

2 CHASSIS TEARDOWN

FIRST UP AS ALWAYS, we try and remove as many of the panels as we can. This just helps keep them safe, stops you from potentially getting scratches on them, and decreases the weight of the whole thing too, making it easier to shuffle and move around when building and installing hardware. The two side panels, once open, can be lifted off their hinges with relative ease, then placed in the chassis box, which comes with some lovely foam packaging. Then once they're out of the way, you can remove the top-most mesh filter, simply by pushing up from underneath.

That's actually a pretty neat design for this kind of case, as with the way the GPU and motherboard are compartmentalized, means hot air should go straight up, pretty much in a chimney style design, which is ideal for efficient heat dissipation. You'll spot that there are a few cables tied down as well, including the front I/O cables. It's worth leaving these in place for the time being while we work on other parts of the build.

3 CPU INSTALLATION

RUMOR HAS IT that AMD intends to go with an LGA socket on its next generation of Ryzen processors, in a similar manner to its Threadripper chips. So, those who are less familiar with Intel's way of doing things might want to brush up on their LGA installation tips. There's no guarantee the AMD socket will look identical to the 6000 series (assuming it's called that), but the likelihood is high it'll be similar to Intel.

To install our Intel Core i5-11600K, simply lift the retention arm located on the right of the socket, then raise the securing bracket up (leave the plastic cover on, you're more likely to damage the pins by removing it before the CPU is in). Then carefully place the processor into position, making sure to align the little gold triangle on the bottom left corner, with the triangle on the bottom left of the socket.

It's almost always indicated on the securing bracket we've just raised, and should also be on the motherboard as well, but if not, make sure the writing on the CPU reads left to right, when the motherboard is positioned like this, with the top left of the rear I/O being the top of the motherboard (in this case, where that 8-pin CPU power is, that's the top of the board).

Once in position, carefully lower the securing bracket, making sure to slot it under the Torx screw below the CPU socket, then carefully reinstall the retention arm back down in place. The plastic cover will pop off. If you can, we recommend putting it to one side (either in the motherboard box or CPU box) in case you need it later.

If you end up with a dud motherboard or processor after troubleshooting, most manufacturers or e-sellers will want you to return the parts for their own testing. Ensuring the CPU socket (for an LGA board) is undamaged and protected from any issues during transit is vital to making sure you observe the warranty and can get a refund or a new board, which is why it's handy to keep the plastic cover somewhere safe.

To re-install the cover, once the CPU is removed, we recommend lifting the bracket up, and carefully installing the cover back on it like that, before securing the bracket back down, so you stay away from the pins as best you can. The bracket itself has little plastic notches that will lock it in position. Again orient the writing so it reads from left to right, in alignment with the top of the motherboard.

6 POWER SUPPLY NEXT? YUP

IT'S AN UNUSUAL STEP FOR US, but before we've even installed the memory or even the cooler, we decided to go with the power supply next. Now, this isn't something we'd traditionally recommend you do this early on in a build, but hopefully, you'll find out just why it's so important to get this done and out of the way as soon as you can.

How the power supply actually installs in the chassis, however, is something of a mystery, unless you have read the manual beforehand. We'd typically expect a case to be intuitive enough through its design that we wouldn't need to do that, hence we didn't open up the manual at first.

Basically, there is a bracket above the motherboard where the power supply sits in place, from which you can route your cables. That then connects to a passthrough power cable which runs to the bottom of the chassis, where you'd plug your power lead into. However, on first inspection, installing the PSU in a traditional manner just looks insanely difficult, as there's no space for a screwdriver in the back of the chassis. The screws that go into the rear of the power supply, are actually facing that rear ventilated area that your AIO is going to be installed in later. Unless you have a tiny, yet hefty screwdriver, it's not going to happen, and even if you did, it'd be fiddly as heck.

What you need to do is remove the entire bracket from the chassis itself, and then install the power supply to that, outside of the case, before reinserting it and securing it back down in the Z5i again. There are five screws that secure the bracket in place, four of which point out towards the side of the case on the motherboard tray side (meaning easy access for our plucky screwdriver), and another one located on the opposite side, near the top of the chassis, in the domain of the GPU. Remove those, and the bracket can then be removed from the case allowing you to install your PSU without worry. Once that's done, you then reinstall that back into the case, and you're good to go.

Looking at the chassis and the lack of cable management options has prompted us to do all of this earlier. We know we're going to be strapping cables to the frame like nobody's business, and trying to route them as best we can. So front I/O cables, RGB, and PSU cables are going in next.

This is one of the odd things about this chassis and goes back to that point we mentioned earlier about visualizing how you think the system is going to be put together. We saw ahead of time that cable management was going to be a nightmare, and so it actively altered how we approached the build. Installing a PSU, threading all the cables through, and handling cable management is usually the final piece of the puzzle. However, in this case, it came pretty much immediately after the motherboard installation because of the situation with the internal cable management. Sad times.

4 M.2 & AIO BRACKET

NEXT UP IS the CPU cooler's bracket for the CPU block, and the M.2. Open your cooler box, identify the Intel bracket, then find the bag with the correct screws for your socket. Pop the backplate into position underneath the board, then screw each of the four standoffs into position. There's likely to be a bit of sag between the standoffs and the backplate, but that's normal.

Now on to the M.2. In the case of MSI's B560i, remove the two Phillips head screws on the "Edge" heatsink, and lift it off. Then insert your M.2 into the port, making sure to align the notches. In the motherboard box, there should be a bag with an additional M.2 screw in it, secure the drive at the other end with that. Once done, peel off the plastic "laird" sticker/cover located on the thermal pad underneath the heatsink, and reinstall the heatsink.

5 MOTHERBOARD IN

HERE WE'VE INSTALLED the motherboard into the Z5i chassis. The rear I/O points downward, as does the GPU, so it is going to be fiddly to install your rear I/O devices later, but that's the price we pay for making it look snazzy.

We've held back on installing the memory, as we're already foreseeing problems with cable management, particularly the cooler. When you're building any system, always take your time to look at how the case is designed, where cables need to be routed, what hardware might get in the way, and plan your installation accordingly.

We've also taken the opportunity to install the PCIe riser cable directly into the motherboard. It's a bit tight but installs exactly like a graphics card. It should already be threaded through that cutout in the bottom, just line it up, and push it into the PCIe slot, until it clicks in place.

7 CABLING HELL

AND SO IT BEGINS.... so here you can see we've begun our cable runs. We're using a multitude of different cable sets to do this. Our short CPU power wouldn't reach the 8-pin EPS on the board, so we've had to go for a long pro cable. We ran a short 24-pin underneath and round through the passthrough to the motherboard and then strapped those down to the aluminum ventilated divider in the case with our own cable ties—G.Skill doesn't provide any with the chassis.

The front I/O is also frustrating to deal with, as it naturally bends out towards the motherboard here. We had to forcefully bend the cables downwards into the "cable management/GPU area" before returning them through the correct passthroughs in the case to get to the correct parts of the motherboard, which was annoying all round.

We also pre-threaded the GPU power and an additional SATA power through round the back, in preparation for installing the 6800XT and our NZXT cooler.

8 BUT THERE'S A PROBLEM...

...AND THE ISSUE IS that our GPU is too tight to the cabling area for us to hide all these cables. The PCIe riser is also secured in position, to allow for three slot cards. So we've detached the PCIe riser cable from the case and let it hang loose (this is how one of our favorite cases, the Hydra Mini, does it). The GPU then sits in those two outside slots, freeing room behind it to tuck all those cables.

There's still a problem with this design. Those cables will insulate the backplate of the card, increasing heat there. The card will also conversely heat up those cables. Lastly, the card's tri-fan cooler design is now far closer to that glass side panel than we'd have liked, meaning higher temps too. So although officially this case supports three-slot cards, we'd recommend sticking to two slot cards.

9 COOLING CONUNDRUMS

DOES THIS LOOK LIKE A MESS YET? As someone who prides themselves on clean builds, it feels like a mess. We've installed the 240mm AIO here, it's a finicky process. We secured the CPU block first, ensuring its digital display is in a legible format, then secured the radiator in position at the rear of the chassis. There's so much wasted space here, and the way it's installed looks a bit haphazard. We'd recommend you go for a full-on 280mm if you can.

Then it was a case of routing the cables from the CPU pump around to that GPU area. We've strapped them to the interior frame again, and run them to the top of where the cables are, again using cable ties. After that, the DDR4 went in, as it was the easiest way to sort the cables out.

10 OH NOOO...

SO YEAH HERE YOU CAN SEE exactly what we mean about the cable management situation. But unfortunately, there is barely anything we can do. We've actually used some of the velcro straps G.Skill has provided to bundle as many cables together as flat as we can, and then strapped those to the SSD mounting locations in the back of the chassis/motherboard tray.

The fan cables and everything cooler-related, and the 8-pin EPS cable, are then strapped to that ventilated mid-plate that separates all three chambers, again using cable ties. We've left the cable tie tails in position for this shot, so you can see exactly where we've done that, but we will cut these off later and hopefully tidy things up a bit.

11 CABLE CALAMITY

AS YOU CAN SEE, things are tight with this cooler. In hindsight, it's not too bad knowing you can rotate the display, but even so, placing the tubing to the right of the block, close to the memory, and the cables to the upper-side of the motherboard would lead to problems with routing as there's so little cable management space here.

Securing the block is a straightforward process. Out of the box, it should come with some thermal paste pre-applied to the bottom. If it doesn't, squeeze a small dollop into the middle of the CPU. Then, making sure you have the correct bracket attached to the CPU block, position the block over the stand-offs, and then secure from each corner using the included thumb-screws. Do this a bit at a time, at each corner, by hand, before finally tightening off with a screwdriver (just make sure it's not electric).

12 THE FINAL PIECE OF THE PUZZLE

AND LAST BUT NOT LEAST we get to that trusty GPU. Admittedly, this was a little more difficult to slot in than usual. Attaching that detached PCIe riser into the bottom can be fiddly but once in place, it's easy to position the card and slot it into the two PCIe covers we removed earlier. Then it was simply a case of securing the card down again with the Phillips head screws, plugging in the twin 8-pin PCIe power cables and we were good to go.

As you can see, this isn't exactly a clean build. There's no space to hide any cables whatsoever. G.Skill could have thought of a solution to this before releasing the chassis. Even a panel between the motherboard and GPU tray area, that you could run cables down and reattach would be handy.

If that means sacrificing some of the case's aesthetics, so be it. Form over function isn't often a recipe for success in something as practical an application as a computer case.

1 In hindsight, going for a bigger radiator, and re-orienting the CPU block in software, would have led to a cleaner build here.

2 Additionally, some decent RGB fans here would help add some light to the interior of this very smoked-out chassis.

3 Custom sleeved cables, cut to a specific length should mitigate the vast majority of your cable management problems, although it'd still be a mess.

4 Again, this case really needs some form of cable management solution in the rear. Wedging them all behind the graphics card just isn't a sensible idea.

THE CONCLUSION

SO, THIS WAS AN experience, to say the least. At first glance, the G.Skill Z5i is a beautiful-looking thing, with curved panels, brushed aluminum, intuitive lighting, and a fantastic premium feel, all for a reasonable price ($200). But the reality is once you get under the covers and start tinkering with the build, the problems just start to spring up.

Are we happy with this build? In a way yes, it fills the brief of building in this chassis, we got it done, in time, with all the hardware we brought to the table without any swaps. But deep down, it's not one we're proud of. It was frustrating to deal with, and the lack of consideration for some of the basics of decent case design resonates quite badly, especially when you do this every day.

"Enough complaining! What would make this case better, Zak?" This chassis is ITX by design, and I know there is a subset of the PC building community that despises large ITX cases, but if I were designing version 2.0, I'd make it slightly taller and wider. Taller to accommodate a 360mm radiator in the rear, and wider to incorporate a smart form of cable management between the GPU and motherboard tray area.

I'd tidy up the internal layout, especially that rear ventilated aluminum area that the radiator sits in. Giving that a cleaner design with obvious mounting locations for the 280/360mm radiator would be top of my agenda. I'd also add an extra 120mm fan mount located on the back of the interior separator that could feed the GPU with cool air pulled from that radiator location, and then let convection channel the heat up and out. Lastly, I'd drop the smoked effect from the windowed panels. Seeing more of the inside would bring attention to both the GPU and the motherboard, and with that cable-management plate in position between the GPU and mobo compartment, you'd worry less about cable clutter showing. Only then would I consider this a 9 out of 10 or higher.

Outside of our mission to get a system built in this, however, there's no doubt that, for a first case, G.Skill is on the right track. As we've said in the past, if you want to break into this market, and capture some sales from the likes of Corsair, Phanteks, NZXT, Fractal, and Lian-Li, you need to bring a flawless product to the table. Right now, the Z5i feels like a concept car. It looks awesome, it has that flying steering wheel, and cameras for wing mirrors, but in reality, it's impractical to use, and that's the problem.

We worry that if this case performs poorly, G.Skill will sack off the division and investment entirely, despite all the hard work that's been put in. It needs a version 2.0, because with some slight tweaks to the design and improvements to the airflow, this thing could be an absolute behemoth of a chassis that takes the fight to the big guns. ⏻

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	2,185	1723 (-21%)
CrystalDisk QD32 Sequential Read (MB/s)	3,522	5003 (42%)
CrystalDisk QD32 Sequential Write (MB/s)	2,988	4,175 (40%)
Total War: Warhammer II (fps)	48	66 (38%)
Shadow of the Tomb Raider (fps)	53	77 (45%)
Red Dead Redemption II (fps)	47	72 (53%)
Assassin's Creed Valhalla (fps)	38	51 (34%)

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

Our zero point consists of an AMD Ryzen 7 3700X, 16GB DDR4 @ 4,133, a Zotac GeForce RTX 3060 Ti 8GB , and a 1TB ADATA XPG Spectrix S40G M.2 PCIe 3.0 SSD. All tests were performed at 4K on the highest graphical profile.

LAB NOTES

IAN EVENDEN, CONTRIBUTING WRITER

Does Windows 365 mean the end of hardware?

Is your home PC about to disappear into the cloud?

MICROSOFT HAS LAUNCHED Windows 365, a way of accessing a Windows PC on any internet-enabled device. It didn't have a particularly smooth launch, as only a day after they began, free two-month trials were withdrawn, presumably because of unexpected demand.

In a nutshell, what Windows 365 means is you rent a Windows 11 PC with a specified number of processor cores, plus RAM and storage, that lives on a server somewhere in Microsoft's basement. You can then access it from your phone, tablet, Chromebook, smart refrigerator, or whatever can run the app required. That poses the question: do you really need a PC in your home any more?

Gamers have already had a taste of this from Xbox Cloud Gaming and Google's Stadia, which are hamstrung by the client's internet connection speed and have some problems that need shaking out. But as home internet connections become faster and more stable, and 5G and its successors fill the air with ever-speedier streams of ones and zeroes, it is possible the day will come when all you need is a phone.

Now, normally, we'd say something here like 'you can pry our graphics cards from our cold dead hands', but we've got years left tinkering with motherboards yet. The cloud revolution is coming (as long as the Sun doesn't destroy all our electronics first) but just as some people have given their music collections over to Spotify, there are those who have doubled down on buying vinyl, and record sales jumped 30 percent in 2020.

The Windows 365 website, available only for businesses and enterprises. No, we're not sure of the difference either.

If that's any guide, then PC building should remain the boutique option, while the majority who don't care about the odd stutter or slow down move over to the new cloud option, and pay a monthly subscription fee. We've been given a vision of the future, and it's not that bad.

ZAK STOREY

Editor-in-Chief

You know what's impressed me lately? A feature on my phone called "Bixby Routines". Now, I can't stand ineffective "smart" AIs such as Bixby, especially the way manufacturers integrate them into their handsets. However, in this case, it's less of an AI feature and more of a standard setting.

What it does is allow you to set certain operating parameters on your Galaxy when different conditions are met. At first, I confined this to wireless charging only. So I made sure once it got to 10pm, fast wireless charging wouldn't engage on my bedside unit, and the fan wouldn't kick in.

But the game-changer was with Waze and Google Maps. I use my phone a lot as a sat-nav, and if you have it set up in its premiere performance mode, with the full 1440 x 3200 resolution at 120Hz, it just chomps through the battery. So with Routines, I have it enabled so that when Waze opens up, it changes the screen resolution to HD+ and the refresh rate drops down to 60Hz. *Voila!* Long-lasting battery life. As I'm not staring at the screen constantly, and it's far away, I don't need it super-sharp, nor do I need it with a high refresh rate. It's awesome. Such a small idea, that just works and works well.

Editor's Pick:
Immerse Gaming Hive Plug-in and Beyerdynamic MMX 300

Can a plug-in really boost your gaming experience?

NOT A DAY GOES BY where I don't listen to music or use headphones in one way or another. Although this probably isn't good for my hearing, audio quality is important to me and I'm always looking to get the best out of the kit that I've come across since working on *Maximum PC*. There are many ways to get the most out of your audio gear and one tech that's been gaining much attention in recent times is spatial audio.

Essentially, spatial audio is AI software that can create a 360° spatial sound field, making the experience more immersive. This added level of immersion is great for gaming, so when Embody's Immersive division offered to let me test its Hive spatial audio gaming software, along with a pair of Beyerdynamic MMX300 headphones, I jumped at the chance.

These are the first Beyerdynamic headphones I've tried and I am impressed. Expectations were high before testing, I've always loved the simplistic looks of the company's headphones and the all-black design with black velvet ear cups on this pair look sleek. They are comfortable too, which is always a blessing, but particularly for a gaming headset. The overall build quality is good, and although the main material for the cups is plastic, it's a soft-touch plastic so feels pretty nice.

First, I tested these bad boys without the plug-in to get to grips with how they sound, and they got off to a strong start. They seem to be engineered for a rich warm sound, whilst managing to not drown out all the detail. I tested a few FPS titles, as well as *Forza* and *Genshin Impact,* and they make a great gaming headset, as well as a decent pair of headphones for listening to music.

A license for the plug-in costs $14.99 for a year and offers a neat experience. The application is easy to use and well designed with clear graphics. During setup, you can choose to send a picture of your ear for a more personalized experience, while on the homepage, there is a visualization of the headphones with three presets: close combat mode (RPGs and racing games); immerse mode for accuracy (FPS games) and; awaken mode (MMORPG and MOBA titles). In-game, you can toggle between the three presets, which is pretty handy.

The Hive experience puts you right in the center of the game and improves awareness when playing FPS titles. I'm not the fastest player—my reactions won't win many esports battles—but having the advantage of 360° audio helps enormously, bringing sounds to the fore instead of being buried in the game's environmental audio. It works with most gaming headsets too, though, headphones such as the Beyerdynamic MMX300 have custom presets. Even after a short time with Hive, I feel it's worth splashing out for the improved gaming experience. **–SL**
Immerse ($15/year) https://embody.co
Headset ($279) https://beyerdynamic.com

The elevated
keyboard marks
the S17 out from
the crowd.

Asus ROG Zephyrus S17
Kooky Keyboard

THE NEW ZEPHYRUS S17 is a beast of a gaming laptop with an interesting USP; a keyboard that tilts up toward the user. At first glance, it's a minor alteration to make typing more comfortable. But while this isn't an inaccurate assessment—the gentle angling of the keys and space to rest your wrists beneath the keyboard make for a pleasant experience—there's more to it than that.

That keyboard lifts away from the chassis to uncover additional vents, allowing the S17 to draw in cool air from the top and bottom of its casing, before expelling it through the rear and sides. A liquid metal cooling compound inside the laptop, combined with two impressively quiet fans, results in a laptop with stellar thermal performance, even with the most demanding games. We wonder about the potential for a build-up of dust and dirt below the keyboard mount, though, which could hinder long-term performance.

The S17 has no trouble handling games, thanks to an RTX 3080 GPU combined with an octa-core i9-11900H processor and a hefty 32GB of high-speed memory. It powers through games of any resolution, making the most of its gorgeous QHD display.

The screen rocks a 165Hz refresh rate and a 3ms response time. Color reproduction is sharp and vivid, and the 17.3-inch panel has a thin bezel for such a beefy laptop. As part of Asus's 'Republic of Gamers' line, the Zephyrus S17 was designed with gaming in mind, so the display is excellent, whether you prefer fast-paced shooters or thoughtful strategy titles. CPU-bound titles excel here thanks to the speedy Core i9-11900H.

We found that 1440p was a comfortable resolution, with few dips below 60fps in even resource-intensive games. Esports titles have no trouble hitting the display's refresh rate; a version of the S17 with a 120Hz 4K screen is also available if you value resolution over framerate. Ray-tracing performance is as good as we'd expect from a laptop RTX 3080, by which we mean it's a definite option, but comes with a pretty heavy performance tradeoff.

One standout element is the SSD, which we don't often say about laptops. It's a 2TB Gen4 NVMe drive, and transfer speeds are nutty; we're talking upwards of 7GB/s sequential reads. Load times are trivially brief, so be prepared to load into a *Rainbow Six Siege* match a full 30 seconds before your squadmates. There's space to add two more NVMe SSDs, and the memory is upgradeable too (the latter is a holdout from the 16GB models, though; 32GB is more than enough).

In terms of aesthetics, the Zephyrus S17 might not be for everyone. The all-metal body is covered in sharp detailing, and the per-key RGB is wonderfully vibrant, red light spilling across the keyboard when you boot the S17 up. It's not exactly lightweight, but we wouldn't expect that from a 17-inch gaming laptop.

The S17 keeps it simple in terms of additional features. A sturdy volume roller rests above the function keys, and a standard 720p webcam nestles in the screen's upper bezel. Connectivity is cutting-edge, with Wi-Fi 6 and Bluetooth 5.2 alongside the physical ports, including high-speed Thunderbolt 4 and the increasingly elusive SD card reader.

Construction is sturdy, with no wobble on the raised keyboard deck. The keys are optical mech switches with an audible click but not much travel. This could be a sticking point for potential buyers; it's not as loud as conventional clicky blue switches but clatters during rapid typing. The trackpad is luxurious, with a firm click and perfect responsiveness.

The S17 is power-hungry, so while the battery does an admirable job, we wouldn't recommend lengthy gaming sessions without a mains plugin. Sound quality is stellar; a six-speaker setup including force-canceling woofers uses Dolby Atmos to excellent effect, among the best we've seen in a gaming laptop.

This is a killer machine, adept for both desktop replacement purposes and on-the-go gaming. The $3000+ asking price (dependent on SKU) is eye-watering but in line with many other laptops. If you have money burning a hole in your pocket and need a new gaming rig, the Zephyrus S17 is a brilliant choice. **–CHRISTIAN GUYTON**

VERDICT

9

Asus ROG Zephyrus S17

➕ **GALE FORCE** Great all-round performance; quiet and cool; excellent monitor.

➖ **GENTLE BREEZE** Fairly expensive; dust build is a concern; keyboard can be noisy

$3,299, rog.asus.com

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	1,030	2,247 (118%)
CrystalDisk QD32 Sequential Read (MB/s)	3,374	7,088 (110%)
CrystalDisk QD32 Sequential Write (MB/s)	2,530	5,189 (105%)
3DMark: Fire Strike (Index)	13,610	24,558 (80%)
Rise of the Tomb Raider (fps)	92	152 (65%)
Total War: Warhammer II (fps)	62	116 (87%)
Tom Clancy's Ghost Recon: Wildlands (fps)	49	85 (73%)

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

Our gaming laptop zero-point is the Acer Predator Triton 500, with an Intel Core i7-8750H, Nvidia GeForce RTX 2060 Max-Q, and 16GB of DDR4-2666. All game tests are performed at 1080p at the highest graphical profile.

SPECIFICATIONS

CPU	Intel Core i9-11900H
Graphics	Nvidia GeForce RTX 3080 16GB
RAM	32GB DDR4 @ 3200MHz
Screen	17.3-inch 1440p IPS screen @ 165Hz 3ms
Ports	1x 3.5mm Combo Audio Jack, 1x HDMI 2.0b, 2x USB 3.2 Gen 2 Type-A, 1x USB 3.2 Gen 2 Type-A, 1x USB 3.2 Gen 2 Type-C, 1x RJ45, 1x Thunderbolt, 1x SD card reader
Connectivity	Wi-Fi 6, Bluetooth 5.2
Weight	5.73lbs
Size	15.5 x 10.4 x 0.8 inches

Acer Swift 3X

Could this be your portable productivity pal?

IT'S FAIR TO SAY that Acer's laptop line is pretty diverse, covering everything from gaming-focused machines to Chromebooks, rugged laptops, and convertibles. It also does an interesting line in ultra-thin laptops, which is where you will find the Swift series. The model we have on review here is the Swift 3X. Just 100 bucks shy of $1000, is this worth considering if you are on the lookout for a new portable work machine?

Despite its 14-inch screen, the Swift 3X is a small and compact laptop. At just 0.7in thick and weighing 3lbs, it's thin and lightweight and easy to take with you on the go. It is light and comfortable to use on your lap, and yet maintains a more rugged feel than you would imagine.

The build quality is good, and while there is nothing groundbreaking about its design, those sleek, minimal looks won't distract from your screen time and tasks. Our model features Acer's steam blue color, which adds a premium-looking finish to the laptop, but it also comes in a sand-like safari gold. The teal blue hinge on the laptop feels a little out of place, but the Acer logos featured on the lid of the laptop and just below the screen are finished in a brass effect and go hand in hand with the steam blue finish.

Opening up the laptop, you are greeted by the 14in full HD IPS display with thin bezels all around. The panel itself provides a good color range and high color accuracy. For creatives and content creators, this is ideal as the screen provides a natural look. However, it isn't the brightest panel out there and you may struggle to use this laptop outdoors.

The 60hz display is suitable for everyday usage, but video editors may wish for a higher refresh rate.

At the top, there's a built-in HD webcam and microphones. Together, these do a decent job when used for video calling, especially considering it shoots in 720p at 30fps. Below the screen is one of the highlights of this laptop, the keyboard. Acer has got just about everything right here—it is comfortable to type on and feels buttery smooth to use for long durations. The backlit keys are nicely spaced out and there is enough room to rest your wrists comfortably on either side of the trackpad. The trackpad itself is responsive and accurate, and though we have seen bigger trackpads on similar-sized laptops, it is good enough for day-to-day use. There is a fingerprint scanner tucked below the arrow keys, a handy feature for a little extra security. It's responsive and a convenient way to unlock the device. In terms of hardware, it has decent enough compatibility, with two USB 3.2 Gen 1 Type-A ports, a USB 3.2 Gen 2 Type-C port that supports Thunderbolt 4, and an HDMI port.

Performance-wise, it features an 11th gen Tiger Lake CPU, the Core i5-1135G7 in this specific model. Along with 8GB of RAM and an Intel Iris Xe Max GPU, which makes it compatible with Windows 11. The Swift 3X can handle tasks simultaneously however, with bigger tasks it could benefit from more RAM. When flicking around the laptop for generic use and web browsing it is pretty fast. Using the Adobe suite or doing light video editing, it is more than capable but struggles with

harder video editing. On the flip side, gaming isn't really the Swift 3X's forte, so if you want this, you will have to look for a laptop with a designated gaming GPU.

Overall, the Swift 3X is an interesting machine. It's extremely portable, and at under $1000, decently priced considering its build quality, productivity capabilities, and hardware. It's a quiet laptop with a decent battery life of up to 14 hours. It's also a joy to use, especially that keyboard. It's just a shame that it doesn't have the horsepower and storage to game and edit videos on, although for gaming at lower settings you will achieve around 60fps performance. If you are in the market for a productivity laptop towards the budget end, the Swift 3X is a good shout, however, there are more powerful products out there in this price range. –SAM LEWIS

VERDICT

8

Acer Swift 3X

◨ **CERTAINLY SWIFT** Great keyboard; easily portable; good color accurate panel; productivity focused; futureproof for Windows 11.

◨ **NOT ENOUGH PUNCHING POWER** Slightly dim screen; not ideal for gaming; needs more RAM; only 512GB of storage.

$899, https://store.acer.com

BENCHMARKS

	ZERO-POINT		
Cinebench R15 Multi (Index)	802	967 (21%)	
CrystalDisk QD32 Sequential Read (MB/s)	3,540	2,302 (-35%)	
CrystalDisk QD32 Sequential Write (MB/s)	2,978	1,102 (-63%)	
3DMark: Fire Strike (Index)	4,849	5,616 (16%)	
Rise of the Tomb Raider (fps)	21	23 (10%)	
Total War: Warhammer II (fps)	31	25 (-19%)	
Tom Clancy's Ghost Recon: Wildlands (fps)	21	24 (14%)	

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

Our gaming laptop zero-point is the Razer Book 13, with an Intel Core i7-1165G7 CPU, Intel Iris Xe Graphics, and 16GB of RAM. All game tests are performed at 1080p at the highest graphical profile.

SPECIFICATIONS

CPU	Intel i5-1135G7 @ 2.40GHz
Graphics	Intel Iris Xe Max
RAM	8GB LPDDR4X @4267MHz
Screen	14.0" Full HD IPS @ 60Hz
Storage	512 GB SSD
Ports	1x Type-C USB3.2 Gen2, 2x Type-A USB3.2 Gen1, HDMI, 1x Mic-in/Headphone-out Combo Jack.
Connectivity	Wi-Fi 6 AX201, Bluetooth 5.0
Weight	3.09lbs
Size	0.9 x 8.4 x 12.7 inches

A sleek ultra-thin aluminum chassis, it looks good on the outside but how does it perform?

ASUS TUF Gaming VG28UQL1A

A 4K panel to satisfy both PC and console gamers?

IF SIZE MATTERS, what's an extra inch worth? Clocking in at 28-inches, the new ASUS TUF Gaming VG28UQL1A is an inch larger than what has arguably become the 27-inch norm for 4K, 144Hz gaming monitors based on IPS panel technology.

That extra inch makes sense, but it's a worry, given that the non-gaming 28-inch 4K monitors we've seen, such as the Philips 288E2UAE and Acer B287K, have underwhelmed in terms of panel quality. Rather than just being a PC gaming panel, Asus has fitted the TUF Gaming VG28UQL1A with two HDMI 2.1 ports. The significance is mainly for console gaming, with both the Sony PS5 and Microsoft Xbox Series X supporting 4K@120Hz refresh rates, which require HDMI 2.1.

Of course, for PC, DisplayPort 1.4 is currently the interface of choice, which enables up to 144Hz refresh at the full 4K. There's a two-port USB 3.0 hub, but no USB Type-C. Asus rates pixel response at 1ms gray-to-gray, while HDR support takes the form of VESA DisplayHDR 400 certification. That's entry-level HDR and doesn't include local dimming. As for color accuracy, Asus claims 90 percent coverage of the DCI-P3 digital cinema gamut, which is decent but not exceptional. Both Nvidia G-Sync and AMD FreeSync adaptive refresh are supported.

In terms of physical build and features, this is a TUF model, rather than Asus's top-notch ROG panels. But quality is decent and the stand is fully adjustable including rotation into portrait mode. But the all-black chassis lacks the snazzier styling elements found on ROG monitors.

All told, it's a pretty sweet 4K gaming package, on paper. And so it proves in practice. Fire up the VG28UQL1A and the first impression of is a vibrant, accurate IPS display with lovely colors, decent contrast, and not too much IPS glow. Forget about those substandard 28-inch 4K monitors, this one looks sweet.

Just as impressive is the speed and response. Asus includes five levels of user-configurable overdrive. The top two suffer from overshoot, enough to be a problem in-game. Levels two and three deliver a sharp response that's up there with the best IPS monitors. Asus's Extreme Low Motion Blur Sync (ELMB) tech is as flawed as ever. It makes little, if any, difference to response while making the display dramatically dimmer.

Refresh rates up to 144Hz are also supported. With 360Hz panels available and 480Hz models incoming, 144Hz is nothing special, but for most people, it will be plenty. Only esports addicts will need more response and latency.

Suitably configured, this panel looks and feels great. Just remember that those uber-sharp 4K visuals make for a pretty outrageous workout for any GPU. You'll need one heck of a graphics card to drive it properly. One other arguable limitation is inevitably HDR performance.

SDR content looks good in HDR mode, so you could plausibly run this monitor in HDR mode all the time. And, given its modest 350-nit rating, it is a bright and punchy panel. Games such as *Cyberpunk 2077* look that bit more punchy and vibrant in HDR mode than SDR mode, which isn't always the case with entry-level HDR monitors. But this still isn't a true HDR monitor.

What it is, however, is a great all-rounder. If you're looking for a high-quality gaming panel to share between PC and console, the ASUS TUF Gaming VG28UQL1A should be on your shortlist. The only catch is that for pure PC gaming, paying extra for those HDMI 2.1 ports is a little pointless. –JEREMY LAIRD

SPECIFICATIONS

Panel size	28-inch
Panel type	IPS
Resolution	3,840 x 2,160
Brightness	350 nits
Colour coverage	90% DCI-P3
Refresh	144Hz
Contrast	1,000:1
Response	1ms GtG
Adaptive refresh	FreeSync, G-Sync
Inputs	2x HDMI 2.1, 2x HDMI 2.0, 1x DisplayPort

HDMI 2.1 will benefit console users, while DisplayPort 1.4 will keep PC gamers happy.

Huawei MateView

Listen carefully, Huawei has a tall tale to tell...

SOMETIMES STANDARDS take hold for no particularly good reason. For proof, look no further than PC monitors and the 16:9 aspect ratio. From a pure computing perspective, there's isn't much reason to favor 16:9 over numerous alternatives. But 16:9 is the standard for the HDTV market, so it dominates the desktop, too.

The exception, of late, has been various flavors of ultrawide panels from 21:9 aspect right out to crazy 32:9 monitors. For productivity, however, ultrawide isn't always super desirable. Enter the new Huawei MateView. It ditches all those wide and wonderful aspect ratios for a much squarer and taller 3:2 format.

Derived from a standard 4K display, the MateView clocks in with 3,840 by 2,560 pixels. So, that's an extra 400 vertical pixels compared with regular 4K. Squeezed into a 28.2-inch panel, it makes for 163 pixels per inch, or around the same pixel density as a conventional 27-inch 4K monitor.

Huawei says the slim-bezel chassis delivers an impressive 94 percent screen-to-body ratio, while the overall style feels distinctly fruit-themed. That's perhaps no surprise given that Huawei's MateBook laptops, which also sport 3:2 aspect screens, pay significant homage to Apple's MacBook portables.

The MateView looks like a strong performer on paper thanks to 98 percent claimed coverage of the DCI-P3 digital cinema color space and true 10-bit color. It's also well-specified when it comes to productivity-related features. Perhaps most significant is the inclusion of USB Type-C connectivity complete with 65W

of device charging. So, you can drive this monitor with a single cable while charging your laptop and connecting peripherals. If you're a notebook user, USB-C is the way to go, it's just so slick and easy. DisplayPort (of the mini variety) and HDMI are also included but note that HDMI is limited to 50Hz.

There's also wireless display mirroring with support for laptops and smartphones, including keyboard and mouse support. The catch? Wireless phone connectivity is only compatible with Huawei handsets and the wireless PC interface is limited to 2K resolution. As for audio, Huawei says the dual integrated speakers serve up theatre-level sound (that's an exaggeration, they are merely adequate), while the dual microphones dampen ambient noise.

But what of the main attraction, image quality? The 4K-plus resolution and 3:2 aspect is certainly great for viewing documents and webpages, with less need for scrolling. With almost 10 million pixels, this is a high-DPI experience, and fonts, in particular, look superb.

It's no surprise the default calibration matches that of a current Apple MacBook pretty closely, but it looks great in Windows, too. Color presets include DCI-P3 and sRGB, but notably not Adobe RGB. This isn't a pro-level content creation panel, but it's still capable of some fairly serious work.

Less impressive is the overall panel punchiness, especially given its 500 nits panel rating. That's partly down to an opaque anti-glare coating that reduces contrast and softens the image quality.

While we prefer some kind of anti-glare coating over glossy-style screens, this Huawei needs a better compromise between reducing reflectivity and maximizing contrast. The pixel response, rated at 8ms, won't win any awards, but it's good enough for a non-gaming remit. HDR is supported thanks to DisplayHDR 400 certification, but it isn't a true HDR display and lacks local backlight dimming.

Ultimately, the appeal of the Huawei MateView hinges on that 3:2 aspect ratio. The panel quality is good, albeit compromised by a poorly chosen anti-glare coating. So, the question is whether the taller aspect suits your workflows. If viewing three or more documents or webpages in parallel is your thing, you'll want something wider. But if more space top and bottom is what you've been craving, the Huawei MateView will be a blessed relief. It's not perfect, but it is a welcome alternative to what has become the widescreen norm. –JEREMY LAIRD

VERDICT

8

Huawei MateView

➕ TALL GUY Intriguing 3:2 aspect ratio; strong feature set.

➖ FALL GUY Poorly chosen anti-glare coating; aspect not ideal for intensive multi-tasking.

$800 consumer.huawei.com

SPECIFICATIONS

Panel size	28.2-inch
Panel type	IPS
Resolution	3,840 x 2,560
Brightness	500 nits
Colour coverage	98% DCI-P3
Refresh	60Hz
Contrast	1,200:1
Response	8ms GtG
Inputs	1x HDMI 2.0, 1x Mini DisplayPort 1.2, USB-C

Huawei's MateView
monitor sports an
unusually tall 3:2 ratio

Glorious curved glass and
integrated RGB lighting make
the G.Skill Z5i a thing of beauty.

G.Skill Z5i
A case of two halves

THE G.SKILL Z5I is one of the best-looking cases we've seen in the last year or so. It would be easy for case manufacturers to fall back on those old tired design tropes, and make something generic. Those age-old cuboid utilitarian shapes are certainly efficient, but it's hardly going to lead to a standout piece of art when it comes to thinking about a home for your next rig.

But the G.Skill Z5i is something else entirely. Its hinged curved tempered glass panels allude to the interior with subtlety and grace. The front I/O is a smooth sheet of brushed aluminum, bound into the rest of the aluminum frame, and littered with triangular perforations that run around and through the heart of the case itself, ensuring that airflow is plentiful. On top is a simple yet effective mesh vent, while the base is a surprisingly premium plastic composite, with integrated underglow RGB lighting, and G.Skill logo cutout, all of which is controllable via your motherboard.

There's no denying that the Z5i is a bold design, the lines are sharp, the aesthetic exceptionally clean, and its premium materials and overall standing look to be near perfect. But does form come before function? Is it just a fraction too much for a company more at home building DRAM than designing the next showstopper?

Don't get us wrong, the Z5i is internally and externally an outstanding piece of modern engineering. Inside, the case has a tri-chamber design. The left-hand side houses the motherboard and your SFX power supply, the right-hand side, the graphics card (and eventually every cable known to man), and at the rear of the unit lies an unfiltered perforated area, where you can mount up to a 280mm AIO, or two 140mm fans along with some 2.5-inch drives. A quick hop over to the Z5i's product page showcases numerous animated videos detailing just how clean your build will look once complete inside. But there's a problem, a big one, and it's something that all those videos carefully sidestep. Cable management.

CROSSED WIRES

The sad fact is there's not a lot you can do here when it comes to that part of our hobby. There are a few cutouts between that center panel separating the right and left compartments, allowing you to route cables from the motherboard to the GPU side of things, and you can hide some cables coming from the power supply itself, but that's about it. You're left to tuck things away around the back on your own.

That means hiding cables behind the graphics card or strapping them to the rear area near the cooler out of sight, and you just end up tying them to the frame in any way you can. This is a case that demands short custom cables at a minimum, and even then, you're still going to end up in a bind. The cable management situation is so difficult to deal with, it completely alters how you build the PC entirely. It's not something you can leave to the end, and tidy up after.

In our build, we ended up cable tying everything down to the main frame that we could, using spare cable ties that we had from previous builds, not ones included with the case itself (although G.Skill does provide you with some slightly annoying velcro-straps). After that, we installed the AIO, before doing some more cable management, before finally finishing up installing the GPU. There's no wiggle room here whatsoever, and that's a shame.

This is one of the best-looking cases we've seen for some time. Its cooling is top-notch (surprisingly), and from a material standpoint, it's fantastic. But, if you can't tick that box of ensuring your PC builder can easily manage their cables, then the design needs to go back to the drawing board. Its functionality lacks the finesse graced upon it by its own form. Sadly, it's that stark contrast between the two elements of this design that makes the Z5i seem so jarring. **–ZAK STOREY**

VERDICT

7

G.Skill Z5i

ZORDON Incredible aesthetic; impressive materials; great concept; good lighting.

LORD ZEDD Cable management; cable management; cable management.

$200 https://www.gskill.com

SPECIFICATIONS

Motherboard Support	ITX
2.5 / 3.5-inch Support	2/3x 2.5 1/0x3.5
Max Radiator Support	240/280mm Rear
Fan Support	2x 120/140mm Rear
Dimensions	12.1 x 7.5 x 16.4 inches
Graphics Card Clearance	12.99 inches
CPU Tower Clearance	2.75 inches

Corsair HS80 Wireless Headset

The most comfortable Corsair headset to date?

AS WELL AS MAKING a great range of PC peripherals, Corsair also prides itself on a solid lineup of gaming headsets across different budget levels. One of its most popular is the HS series and the latest addition to its arsenal, the HS80, takes a more premium approach to the rest of the bunch. So, is it a headset that should grab your attention?

The high-end approach is noticeable straight out the box. More premium materials have been used, including machined aluminum and a soft-touch rubberized plastic on the headband and in the center of the earcups. The only parts of the headset that feel cheap are the cups—it would have been nice had they been made out of the same soft-touch rubber found in the centerpieces.

But on the whole, it feels a step above previous HS series models. Here, everything feels nicely weighted, including the swivel and pivot movement on the earcups, which gives you confidence in the quality. The headset feels durable, as it should do considering the $200 price tag. The design takes more inspiration from Corsair's Void series than its HS series and looks pretty slick, with an all-black color scheme and small iCUE RGB logos on each earcup. The cans are fairly large, but the soft fabric material cushions the ears.

The suspended elastic headband is one of the key differences between the HS80 and other Corsair headsets. We saw this feature on the Steelseries Arctis 9 headset we reviewed recently and as it takes the weight off your head, you can wear these cans for longer. The elasticated band also has a soft leather padding to cushion the head even more so, you can use this headset for hours without them causing any strain. They aren't the most comfortable headphones we've tested, but they are more than good enough for a long gaming session.

The HS80s have some neat features too. On the left can, there is a foldable microphone made of rubber. It's not as flexible as bendy metal microphones but stays firmly in place. There is also an RGB indicator on the end that glows white when you are live and red if you are muted while raising it upright quickly mutes the sound. The microphone quality provides clarity through voice chat and party modes, even on a busy server. There is minimal delay and as far as gaming goes, it does a great job. On the left ear cup, you will also find the volume wheel and power button, which are easy to locate and press on the fly.

MISSED CONNECTIONS

One issue with the headphones is the limited connectivity options. Bluetooth and auxiliary are out of the equation, but Corsair makes up for this with its compatibility. The HS80s work with PC, Mac, PlayStation 4, and PlayStation 5. You can connect these headphones through the heavy-duty braided USB C cable provided, or wirelessly with Corsair's hyper-fast Slipstream connectivity. The wireless connection pumps out high-fidelity sound with ultra-low latency wireless speeds and a pretty impressive audio resolution of 24-bit/48kHz. It also boasts up to 60ft of wireless range and up to 20 hours battery life. However, for true high-fidelity sound, your best option is to go the wired route with the USB C to achieve a 24-bit/96kHz sound.

So, how do the HS80s actually sound? The detail delivered by these cans is superb. Directional audio is pinpoint, putting you right in the middle of a game and giving you that full immersion. The HS80's stunning detail keeps you alert and one step ahead of the opposition. It is a well-balanced sound, bassy enough without being overwhelming.

Though the headphones are optimized for gaming, music consumption shouldn't be overlooked and they deliver a good punchy performance. PC users can take advantage of Dolby Atmos spatial audio to improve the headset's performance. Using the Dolby Access app is an easy way to individually tweak the EQ to your liking, provided you don't mind paying extra for the service.

The HS80s are the leading model in the H series and are currently one of the best Corsair headsets on the market. The gaming performance is highly immersive and the clarity is superb. Although it doesn't have all of the connectivity we would like, it makes up for this with strong compatibility across a range of formats, great comfort, and good build quality. It puts up a hell of a fight for itself in the gaming market with its performance and microphone detail. However, just a few minor details are holding it from getting the full kick-ass verdict. **–SAM LEWIS**

VERDICT
9

Corsair HS80 RGB Wireless Headset

➕ **BATTLE READY** Great accurate clarity; excellent microphone; durable build; good comfort.

➖ **LEFT THE BACKPACK BEHIND** Few cheap materials used; no Bluetooth or auxiliary support.

$199.99, www.corsair.com

SPECIFICATIONS

Driver Type	Custom 50mm neodymium driver
Frequency Response	20Hz-40,000Hz
Impedance	32 Ohms
Connectivity	USB C, Bluetooth
Compatibility	PC, Mac, PS4 & PS5
Battery life	Up to 20 hours
Design Style	Closed back

Asus ROG Falchion

You Have My Sword

ACCORDING TO WIKIPEDIA, the free encyclopedia, a falchion is a 'one-handed, single-edged sword of European origin'. It's clear how this factors into the naming convention at play here, since Asus also offers the chunkier Claymore keyboard within its Republic of Gamers brand, which can swap between 100 percent and 80 percent configurations. We're half expecting it to release the 'Zweihander', a keyboard that somehow boasts a mythical 120 percent layout.

So, the Falchion is a compact gaming keyboard, with a 65 percent layout (making it slightly bigger than conventional 60 percent keyboards) and wireless connectivity. It has all the usual trappings of a premium gaming keyboard; per-key RGB lighting, Cherry MX mechanical switches, programmable macro keys. Straight out of the box, this feels like a premium product. The price tag matches, too; $160 is no small order for a gaming keyboard, but the Falchion does the work to justify the cost.

The overall build quality is great, with an aluminum top plate and double-shot textured PBT keycaps that feel satisfyingly grippy. On the underside of the keyboard are four rubber pads, with folding feet at the rear so you can angle the keyboard toward you. The Cherry MX Red switches in our review model have very little wobble and feel good for both typing and gaming; MX Blue, Brown, and Silver Speed models are also available.

The Falchion can be used in both wired and wireless modes, with the former charging the keyboard in the process. Wireless is achieved via a 2.4GHz USB-A dongle (housed in a nifty magnetic slot when not in use) and offers near-perfect input with no noticeable latency. The battery life is reportedly up to a phenomenal 450 hours, and while we weren't able to verify this, we did find that with RGB lighting at full brightness, the Falchion operated for over 50 hours without needing a charge. Battery life is indicated by a small light bar built into the left-hand side of the chassis.

Below this light bar is the Falchion's best innovation; a touch-operated volume slider, which lets you adjust the volume with quick taps or controlled swipes of your fingertip. It's a great inclusion, adding dedicated volume control without demanding more space on the frame like a roller or extra buttons would. The light bar shifts to indicate the current volume when the slider is touched.

That 65 percent configuration is pretty sound, even for those who prefer full-scale keyboards. Asus has wisely opted to include the arrow keys here, as well as Ins, Del, PgUp, and PgDn squeezed in on the far right of the keyboard. Secondary key functions are clearly labeled in white on the front face of each keycap, with the arrow keys working in conjunction with Fn to swap between lighting profiles and brightness on the fly.

Speaking of lighting profiles; ROG Armory Crate is present for adjusting RGB effects and reprogramming keys. Sadly, it remains a second-rate piece of software that hands off some of its responsibilities to a second program (Asus Aura Creator). Armory Crate is functional and has a better UI than some proprietary peripheral software, but it's still far from perfect.

The one slightly baffling part of this product is the polycarbonate cover that comes with it. This is essentially a plastic rectangle that covers the keys when not in use, or can be used as a mount of sorts for the entire keyboard. The Falchion is theoretically a great on-the-go keyboard, so including a case of some description feels warranted. However, this cover doesn't properly attach to the keyboard (it just rests on top) so it feels less protective than a proper case and, at worst, a cumbersome addition.

It's fair to say that our gripes are small, though. The ROG Falchion is one of the best compact keyboards we've seen, between its durable construction and solid typing experience. The RGB lighting is bright and colorful, and the volume touchpad is an inspired addition to an already good peripheral. If you're after a new 65 percent keyboard, this is the one. –CHRISTIAN GUYTON

VERDICT

9

Asus ROG Falchion

➕ **SHARP EDGE** Great battery life; smart 65 percent key layout; sleek volume slider.

➖ **DULL BLADE** Keyboard cover is a bit rubbish; Armory Crate remains sub-par; a bit pricey.

$160, rog.asus.com

SPECIFICATIONS

Switch Type	Cherry MX RGB Linear Red/Clicky Blue/Tactile Brown/Silver Speed
Switch Durability	50 million presses
Media Keys	Integrated alt-keys
Connection	USB-A
Battery Life	450 hours maximum
Wrist-rest	N/A
Dimensions	12 x 4.0 x 1.5in

Asus ROG Gladius III Wireless

Asus cuts the cord but the Gladius is still held back

ASUS'S ROG GLADIUS gaming mice have been kicking about for over half a decade now, with each subsequent iteration improving slightly—and only slightly—on the previous model. Unsurprisingly, this is still the case here, with the Gladius III Wireless removing the power cable but ultimately doing little to differentiate itself from the standard Gladius III.

Is that such a bad thing? This is still a well-made device that ticks all the standard boxes for a gaming mouse. It features three-zone RGB lighting (customizable via the Armory Crate software, or the lighting profile button on the underside) along with a sturdy construction and swappable Omron switches for the two main clicker buttons thanks to Asus's push-socket design. This allows you to switch over to quieter switches, or merely extend the mouse's lifespan with replacement parts.

The default Omron switches offer a firm click, as do the two thumb buttons on the right-hand side. Sorry, southpaws; this is a right-handed mouse with no options for physical customization. The scroll wheel is translucent with a rubberized finish, which feels great for controlled scrolling. A dpi adjustment button sits behind the wheel, well-positioned to avoid accidental clicks.

The exterior of the Gladius III Wireless is exclusively black plastic, with a glossy finish and texturing on the sides to improve grip. It's quite tall for a gaming mouse but has a nice arched form that makes it comfortable and ergonomic. Despite losing some of the weight of the original Gladius III, it's still a bit heavier than we'd like, and the weight is noticeably distributed towards the rear. It's not ideal for intense games that demand quick mouse movements, especially if you prefer a fingertip grip.

Connectivity is tri-mode; Bluetooth, 2.4GHz, or wired. None of these modes generate any appreciable input latency, although we found the 2.4GHz more reliable than the Bluetooth. You need a USB-A port on your device for the included dongle, but the mouse has a slot for housing it on the underside so it's easy to take out and about. The plug-and-play functionality of the 2.4GHz mode is perfect for pairing with a gaming laptop.

In terms of general performance, the Gladius III Wireless does well. Battery life is great, with moderately better performance using the Bluetooth mode or by turning off the RGB lighting. A full charge lasted six days (about 48 hours of use) in 2.4GHz mode with the lighting on, which is certainly above average.

The optical sensor is a 19,000 dpi model, tuned up to 26,000 by Asus. In practical terms, the average gamer is unlikely to want much above 3,000, but high-dpi sensors have become the norm for gaming mice and this Gladius doesn't buck the trend. Lift-off distance is good, with input feeling snappy and responsive. Teflon feet on the base let the mouse glide easily over any surface, and a set of replacement feet are included should they become worn or damaged. Between that and the extra Omron switches, this feels like a long-life product.

The Armory Crate software isn't the best, but it can be used to reprogram the six available buttons. The RGB lighting can be customized too, with the scroll wheel, ROG logo, and left-side detailing illuminated separately. The Gladius III Wireless looks good, with a decidedly minimalist design compared to many other gaming mice.

Our biggest issue with this Gladius is that for the price, it doesn't set itself apart from the crowded gaming mouse market. That's not to say it's a bad product, but there are plenty of competitors in this price bracket that do more interesting things. If you're a palm-grip user looking for an ergonomic mouse for work and play, the Gladius III Wireless is a good shout. Otherwise, we recommend you do some comparisons to find the best mouse for you. **–CHRISTIAN GUYTON**

VERDICT

7

Asus ROG Gladius III Wireless

➕ **GLADIATOR** Good connectivity options; swappable switches; excellent battery life.

➖ **CAESAR** Poor weight distribution; lacks unique features.

$90, rog.asus.com

SPECIFICATIONS

Sensor	PixArt PAW3370
Sensitivity	19,000 dpi (tuned up to 26,000)
Sensor Model	400 ips
Polling Rate	1000Hz
Programmable Buttons	6
LEDs	3 RGB zones
Cable Length	200cm
Weight	3.2 oz

Elgato Facecam

A streamer's right-hand man?

WEBCAMS HAVE BECOME an essential part of day-to-day life for millions of people over the past year or so, but many are now ditching their built-in cams for more premium peripherals. The Facecam is Elgato's first webcam, so should people turn their attention towards this cam and add it to their setup?

Considering the huge role it plays in streaming peripherals, it's odd that this is Elgato's first go at a physical webcam. Sure, the company released the EpocCam, which turns your phone into a webcam via an app, but it has now opened the door into webcam production. Hopefully, this camera can complete Elgato's arsenal of streaming gear.

Straight out the box, we have to say this is a chunky cam. This isn't necessarily the worst thing as it houses a pretty big-looking lens, which is the most important part of a camera, right? Included with the Facecam is a little lens cover which is always a nice addition to the package. Nothing says privacy like physically blocking out the camera. Also, the cam comes with a detachable L mount that is quarter-inch thread compatible. This pairs wonderfully with Elgato's Master Mount System too if mounting it on top of your monitor isn't the best setup for you.

However, the standard mount isn't the greatest and can feel a little weak compared with rivals, such as the mount on the Razer Kiyo Pro that we tested in a previous issue. Another problem is that you can't adjust the horizontal angle of the Facecam when it is mounted, as you will either tighten it or loosen it.

On the whole, the Facecam feels a little cheap due to its all-plastic build and is pretty light. However, the lightweight nature will put less strain on your monitor so every cloud has a silver lining. One final point to mention is that this webcam connects with USB 3.0 and comes with a cable too.

Setup is a simple process, and while you can just plug and play, we would recommend getting Elgato's Camera Hub application to tweak some of the camera's settings. Conveniently, as soon as we plugged it in, the video quality looked great. Considering the testing environment wasn't particularly well-lit, the Facecam did a great job at maintaining crystal clear clarity and making sure the image was bright but not over-exposed.

The lens used in the Facecam is Elgatos prime F2.4 lens, with a 22mm focal length. It also features Sony's STARVIS CMOS sensor, which is pretty good going. The 82° field of view will cover most of your room and because the Facecam is limited to Full HD 1080p, unfortunately, it isn't a 4K cam. However, if you use the Facecam mainly for video calling, that should still be sharp enough.

If your image isn't looking quite right, the Camera Hub app is a decent place to fix some settings. Thankfully, like the Facecam, the app is easy to set up and simple to understand. Once inside the app, you will be greeted by a preview panel and the tools on the left-hand side of the screen. These are all pretty generic settings, but it makes it easy to tweak different elements and get your image looking its best.

You can also undo settings individually or factory reset the camera. Settings are also saved directly on the webcam so you can move it between different setups. One thing to note is that the lens is fixed focus, so in the app there is no way to tweak this, although your whole shot should be in focus.

For its first attempt at a physical webcam, Elgato has done a grand job. The main thing to point out is that the image quality is clear and looks buttery smooth at 60fps even if the lighting isn't the best. With no compression and low latency, it makes for a detailed image. Unfortunately, there is no HDR which is a shame, however, colors still look deep and you can alter the white balance and saturation settings in the Camera Hub.

Other negatives worth mentioning are that the build quality and design aren't the greatest, there is no built-in microphone, and the mount isn't the most secure. But if you can look past this and the slightly pricey cost, the Elgato Facecam is a capable camera to add to your setup. **–SAM LEWIS**

SPECIFICATIONS

Video resolution	1080p @ 60/30 FPS, 720p @ 60/30 FPS, 540p @ 60/30 FPS
Sensor	Sony STARVIS CMOS
Field of view	82°
Focus type	Fixed
Connectivity	USB 3.0 A to C
Weight	0.21lbs

Play co-op mode, with three pals and a jellyfish

ACTION ROLE PLAYING

The Ascent

C-beams glitter in the dark

THE INTRO TO BLADE RUNNER, but with way more shooting, seems to be the elevator pitch here, and it works extremely well. As an indentured laborer on the planet Veres, paying off the cost of your new life on the off-world colony with your body, you are at the bottom of the pile. Quite literally. They call it the Deepstink, and you're sent down on a new-character-in-an-RPG mission to reset the computer that's not clearing the sewage so your betters, who live higher up the arcology, aren't showering in manure. You need somewhere to ascend from, after all.

LOCKDOWN MISSION

It would be hilarious if it wasn't so dangerous. This early mission is straightforward enough—you go in with your gun, blast a few ferals, hack a few doors, watch some robots at work, reset the computer's memory. It was a bit like getting back into the *Maximum PC* office after Covid lockdown was lifted.

Shooting is a twin-stick affair, the left one handling character movement while the right controls shooting direction. The camera is fixed, which can occasionally mean some higher part of the level getting in the way, but rarely. There's a dodge-roll to get you out of packs of melee-

only enemies and dodge the telegraphed incoming artillery fire, but most of the time you'll just have that trigger clamped down as you dance your laser sight across the pack that's running at you.

There's a cover system, mastery of which is key to combat, but early opportunities are few. So many fights turn into a procession, running backward while kiting a few ferals that mindlessly run at you. The first time you encounter a boss it's in an open donut-shaped arena with few places to hide. The cover does come, as pieces of the disintegrating superstructure come to your aid.

The ability to aim high, shooting above the heads of smaller enemies and over cover too, comes in handy with larger bad guys, especially as it has a chance to stagger them, and lurking at the bottom of a staircase is a great way to set up a killing ground. The game can be played solo, or in co-op with three others, including a wonderful couch co-op mode which is perfect for the game. Combat is fun, but not exactly subtle, largely becoming a case of how much damage you can send the enemies' way while mitigating the effects of what comes back the other way.

The hyper-urban futuristic setting has more than a whiff of Games Workshop's

Necromunda about it, as well as the neon-in-the-darkness of *Blade Runner*. There's some William Gibson too, and just about any 1980s sci-fi that imagined a high-tech future that still used CRT screens. It's hugely detailed, with every level consisting of a critical path and an absolute mountain of stuff that's not on it, but exists to draw your eye. Robots weld in the dark, lighting themselves (and often more interesting things too) with showers of intermittent sparks. Bodies lay huddled at the edges of walkways, the 'Investigate' prompt unleashing a kick that rattles currency free.

GETTING YOUR KICKS

Trampling all over boxes and crates breaks them to reveal loot left inside, exposing a playful physics engine that sees you knock chairs flying and kick the bodies of enemies a surprisingly long way after you've downed them. Health drops are common enough, at least in the beginning before the cover has asserted itself, that pray-and-spray tactics are perfectly suitable, especially once you've picked up a fully automatic weapon, which doesn't take long. Being an action RPG, there are few places you go that won't dissolve into a fight of some sort, the
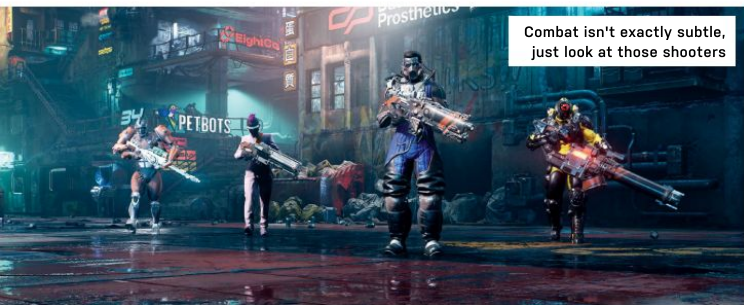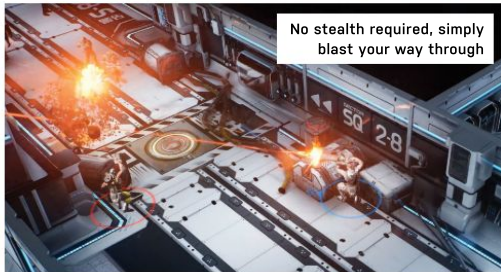
Combat isn't exactly subtle, just look at those shooters
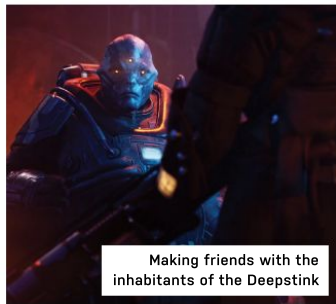
Making friends with the inhabitants of the Deepstink

Visit the Node bar during downtime between scraps

No stealth required, simply blast your way through

Gritty nocturnality, with a bit of 80s sci-fi thrown in

loot you scoop up coming in handy during downtime between scraps when you can upgrade everything from your gun to yourself. Leveling up gains you skill points, which can be spent to improve your weapon handling and maximum health, among other attributes.

Running in 4K with ray-tracing enabled on our RTX 3080, this contrast of darkness and bursts of lighting, often in motion thanks to the numerous fans bringing fresh air to the fetid, can be a remarkable-looking place to explore. When not covered in trash, the floor shines a dull metallic gray, and anywhere with shop fronts demands you to keep still and take it all in. It's a shame that you have to keep moving, as it's in the stillness that the extent of the environmental art and imagination makes itself known. Later levels, in cleaner, better-lit parts of the towering city, are almost an anticlimax

following the gritty nocturnality of the game's opening.

Your ascent from this sewage-clogged hell depends upon you becoming a hired gun for the gang leaders and other characters you encounter. It's here that *The Ascent*'s weaknesses start to show themselves. The writing is by turns sophomore sniggering and banal grunting that has you reaching for the skip button. Then there's the unchanging grimness of everything—even as you ascend the tower, the stench of corruption never really leaves, it just changes form. Add to this the fact that shooting is your only way through, there's no way to use the skill tree to create a stealth character or a hacker, and the whole artifice can begin to feel a little hollow, like a set that might come crashing down at any moment.

The sparks, neon, and laser rangefinders of *The Ascent* may be

spectacular in contrast to the darkness of the opening levels, but that spicy heterogeneity, along with the solid shooting, isn't quite as compelling as the game moves into its middle section. It's fun in single-player mode and a blast in co-op, but this cyberpunk fantasy doesn't quite retain your interest all the way to its beautifully lit elevators. **–IAN EVENDEN**
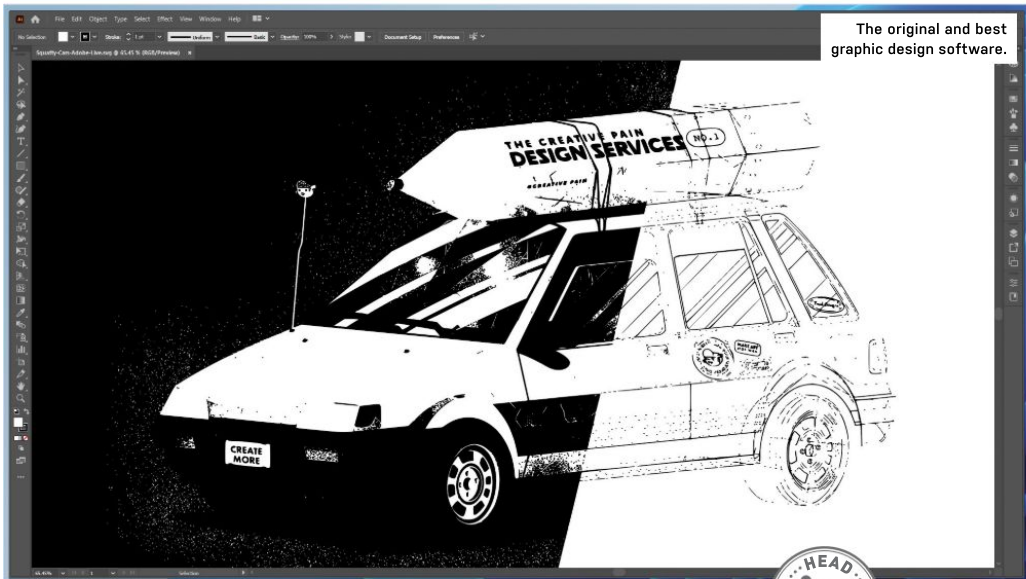
The original and best graphic design software.

Adobe Illustrator vs. Affinity Designer

Can the design newbie knock Adobe off its perch?

BEFORE WE PIT THESE two rivals against each other, let's take a step back in time to where it all began. The year is 1987, Adobe has just released its first piece of software called Illustrator for the Apple Macintosh. It was, and still is, a vector art and design software that by today's standards looks pretty darn basic. However, in the late 1980s, this piece of technology completely revolutionized the art of graphic design.

Before computerized work, creating a piece of design was a massive task and extremely time-consuming. Everything had to be done manually, including rulers and french curls, and typography was created letter by letter. One page of set type may even take a day or two to complete. You think making a mistake on a computer is a problem? Nowadays, you simply press CTRL-Z to undo the error, but back then, a mistake in traditional design work could often mean starting all over again.

Illustrator was designed to prevent this from happening while being more efficient and pushing the boundaries of what could be done. The concept of a computer being able to draw for you introduced a whole new way of thinking to industries such as graphic design, marketing, and publishing. Illustrator eventually replaced Adobe's PostScript software, a page description language that was previously only usable by computer programmers. Illustrator was an application that made the process a whole lot smoother, allowing more accurate control over curves, shapes, and lines that you could see in front of you.

The pen tool was one of the first and most useful tools Illustrator had. It meant you could draw accurate curves, perfectly straight lines and also create infinite variations of shapes by drawing as many points as you wanted. These were the building blocks to every type of design you see today and that certainly gives Illustrator some brownie points.

There have been lots of updates and new versions of Adobe's Illustrator since, but its purpose has never changed. It is engineered to be a graphic designer's best tool, used to create brilliant vector designs that allow you to make graphics, logos, icons, or typefaces that you can scale up and down without losing resolution. Illustrator also comes with specialized tools that help you create more complex designs to perfect your work.

Templates are another great feature, removing the need to create documents from scratch. There are templates for digital, print, web, and mobile, that support both CMYK and RGB color spaces. Illustrator's custom workspaces can help you organize the UI exactly how you want it—that's the great thing about Illustrator, it's extremely adaptable. It also works with the rest of the Adobe Suite of programs, including Photoshop and InDesign.

But now we have a contender, in the shape of Affinity Designer. Serif also began in 1987 creating creative software for Windows, but later launched the Affinity range of applications, the first of which was Affinity Designer, released in 2014. This application is aiming to set a new industry standard for graphic design, and with some decent awards in recent years, including winning Application Creator of

© ADOBE
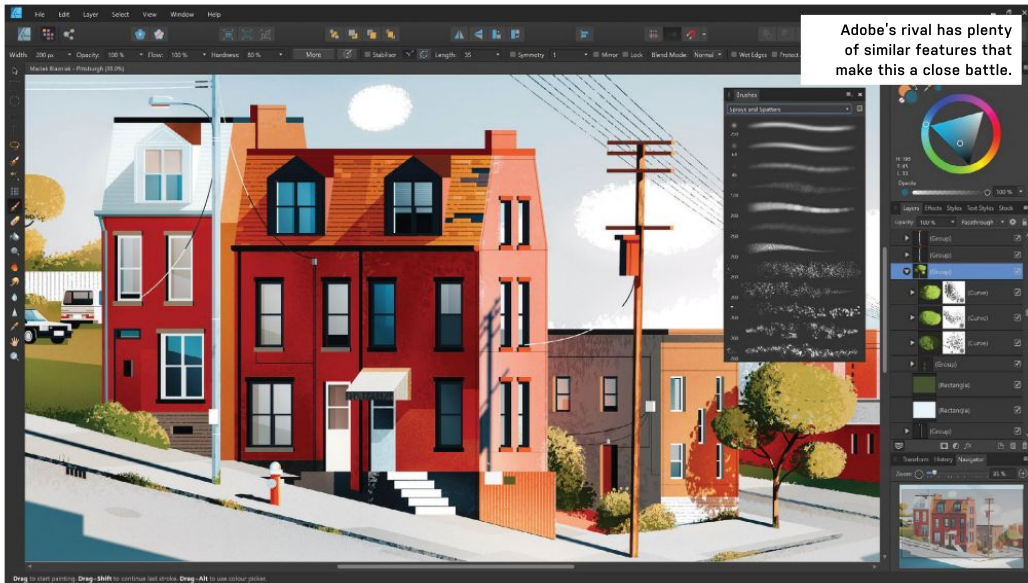
the Year in the Windows Developer Awards 2018 and a five-star award from Editors Choice, it has become a real alternative to its illustrious rival.

Affinity Designer is a similar program to Illustrator—both applications offer the user a similar experience and both are used for the same end products. However, there are some areas where Affinity Designer beats Illustrator. One of the main benefits of Affinity is the lack of a subscription fee. A one-off payment of $49.99 is far cheaper than paying a minimum of $20.99 a month for Adobe Illustrator on an annual plan.

Illustrator hasn't always been subscription-based, however, it was brought under Adobe's Creative Cloud subscription service, which now costs $79.49 per month for all of its programs. If you don't want to sign up for an annual contract, that's some pretty big bucks. Affinity takes the crown here on affordability, however, both programs offer a trial service, so we'd recommend trying them before committing to a purchase.

Affinity is also a marginally easier piece of software to use. Its layout is cleaner, bolder, and more colorful, whereas Illustrator can be a little overwhelming, particularly for a new user. If you have used other Adobe products before, Illustrator's layout should be familiar, but approaching both programs as a new user, Affinity's product has a more organized layout by default. It's also a lot more fluid and visually speaking, benefits from running at

60fps. Due to this ease of use, Affinity could be more appealing to a newer audience, however, it doesn't have as many tools and features as Adobe's offering.

On the whole, both of these programs are excellent, offering great services to creative types, whether they use these programs for a hobby or take a more professional interest. Illustrator lends itself slightly more to professional usage. This comes down to its extensive history of use within the publishing and design industries, meaning that more people are accustomed to using this application.

Affinity Designer may appeal more to the beginner with an easier user experience and it opens the door for newer creatives to work digitally. With a cleaner and faster UI, it makes Illustrator look heavy and will certainly gain more traction because of this. Also, Affinity's subscription-less program gives it an edge over Illustrator when it comes to fresh artists and creatives looking to save a substantial amount of money per year.

They both have great qualities but in terms of professional usage and more features, Illustrator is worth that extra bit of patience to learn and use as it also ties in well with other Adobe programs. For an overall package and longevity purposes, Illustrator has the edge over Affinity design due to its extensive background and business hierarchy. Nevertheless, that's not to say that Affinity isn't an excellent product too. –SAM LEWIS

VERDICT 9

Adobe Illustrator

➕ **INDUSTRY STANDARD** Feature rich; revolutionized graphic design; powerful program.

➖ **STANDARD INDUSTRY ISSUES** Subscription based; slightly more cluttered interface.

⚙ **RECOMMENDED SPECS** Operating system, Windows 10 (64-bit) versions V1809, V1903, V1909, and V2004. CPU, Multicore Intel processor (with 64-bit support) or AMD Athlon 64 processor. RAM, 16 GB. GPU, OpenGL 4.x.

From $20.99 per month www.adobe.com

VERDICT 9

Affinity Designer

➕ **A THREAT TO THE TOP** One-off purchase; easier to use; fast and responsive @60fps; vector and raster.

➖ **CLOSE, BUT NO CIGAR** Doesn't quite have the same history; still not industry standard; fewer features.

⚙ **RECOMMENDED SPECS** Operating system, Windows 10 (1607 Anniversary update or higher), Windows 8.1, Windows 7 (Service Pack 1; Platform Update KB2670838; Aero enabled). CPU, N/A. RAM, 4GB RAM. GPU, DirectX 10-compatible Graphics Cards and above.

$49.99 www.affinity.serif.com

LETTERS

> Folding@Home
> Big Brother
> Windows 11 Compatibility

What PCs are all about

I am a digital subscriber in the UK, and have enjoyed reading *Maximum PC* over the past few months. I am old enough to have experienced the joy when massive hard disks first became available—I refer to 40MB hard disks, of course. Not too long after that, I was responsible for buying a server running a version of Unix, that had two enormous hard disks, that would handle the company's storage requirements pretty much forever. One of these was 500MB and the other was 1GB. And the 1GB disk cost more than my house. *Tempus fugit*, as they used to say back in Rome. I am confident that I will have 100TB of storage in my home computing device at some point in the not too distant future. I don't know what, but I am confident that something will want the increased acreage.

However, the real purpose of this mail was that article on Folding. This was immense! This is what computing is all about—solving problems with computers' ability to crunch data, and this Covid problem is very real and present to ALL of us. If I had a print edition, it would have been circulating among my friends but, unfortunately, the only version I have access to is on my Kindle Fire. Is there an online version of that article anywhere that I could point my friends to?

–P. Revell

EDITOR-IN-CHIEF, ZAK STOREY, RESPONDS: Hi there, and welcome to the *Maximum PC* community. We actually have a fair few readers from the UK, at least as far as I can tell from our internal stats and databases, and you're more than welcome here.

You're absolutely right, I was blown away by Ian's piece on Folding@Home and how it all came together. His work, particularly in long-form journalism, has been outstanding over the last year, and something we're very proud of. Every time his work comes in, it's captivating to read, especially with what he covers, and how he covers it. His work on the Colossus machine, the first electronic computer, developed and used in WW2 to crack the Lorenz SZ cipher back in our May issue, is also well worth checking out. And of course, he's also written a phenomenal piece on the Y2K bug in the March 2021 edition. In fact, he's impressed me so much, we've scheduled a feature from him in each and every issue (including this one) up until early next year, and likely beyond that too. Fantastic stuff all round.

Sadly there's nowhere online you can grab this right now. I have spoken to the senior management at *Future* about setting up the PDF archive online again, however, they're reluctant to do so—we'll let you know if anything changes. Until then though, all we can do is recommend that you keep reading your digital subscription.

And you're right, this is exactly what the PC is all about. We obviously spend a lot of time here at *Maximum PC* talking about gaming, and the latest graphics cards, and how they work, and also a significant chunk of the issue building PCs. But without a doubt, it's these kinds of projects and applications that will push us further as a species.

Big Brother

Hello Zak, a loyal, Canadian reader here, as your magazine is most readable and understandable. In the Build It section, August 2021, "A Pseudo Retro Conversion", the cost of ingredients is $1,984. How Orwellian! In this build, you took us from Functional Funk, to Magic Time, to Riser Wizardry, and finally to no monitor. Sadly, it would not have been a CRT. This took me back to earlier micro-computing days and I thank you very much: Orwell and No Stock issues aside. As you say this isn't a great time to be computer-building and here in Canada the cost is truly out of reach for most builders.

In this Covid-ending time, fingers crossed, please keep up the great work. Let's hope prices come down to sane levels so that we can employ Build It to experience what's new in the PC world.

– Murray

EDITOR-IN-CHIEF, ZAK STOREY, RESPONDS: Oh my days, how did I miss this? I love a good pun, like seriously. Bad puns, good puns, doesn't matter!

↘ submit your questions to: editor@maximumpc.com

But yes, that was entirely an accident and something I'll never be able to repeat again. As a man who spent far too much of his teenage years reading *Animal Farm* (thanks weird educational curriculum), and getting way too bored of Communist piggies, I'm sad I didn't spot that, otherwise I would have gone overboard on every 1984 pun I could muster.

Anyway, we are slowly seeing a decline in costs for graphics cards now, and stock does seem to be eking its way back into availability, albeit slowly. Issue on issue we're seeing improvements, so that's all good. As the world recovers from the global shutdown, and working procedures change and new fabs come online, I imagine we'll start seeing normal availability from early 2022 once again. No doubt spurred on by China's crackdown on crypto as well. That might quell the potential for Chia coin SSD shortages too, thinking about it. Only time will tell.

Windows 11 Compatibility

Dear Zak, I'm a long-time subscriber to *Maximum PC* and though I am not a gamer, I look forward to each issue. The main reason for this communication has to do with Windows 11. As we all know, on June 24th, Microsoft made its announcement for Windows 11. Along with that, it released the PC Health Check tool as a download.

To the astonishment of most of those that downloaded the tool, they discovered that their PCs were not suitable for Windows 11, even though they are currently running Windows 10 without issue.

Industry watchers and leading publications quickly picked up on the consternation generated by machines running Windows 10 suddenly being "incompatible" for

Windows 11. As we also know, Microsoft got heat and actually admitted that the PC Health Check tool has "bugs", for which they then provided an updated version. But shortly thereafter Microsoft actually pulled it from general availability because of other "issues".

The technical requirements for Windows 11, as stated by Microsoft are clear enough: Processor: 1 GHz or faster, two or more cores on a compatible 64-bit processor or system on a chip (SoC), RAML 4 GB or greater, Storage, 64 GB or more, Graphics Card: Compatible with DirectX 12 or later, with WDDM 2.0 driver, System Firmware: UEFI, Secure Boot capable, TPM: Trusted Platform Module 2.0, Display: High Definition 720p @ 8-bit/channel, Internet connection, Microsoft Account to complete setup.

The rub comes with a) UEFI, b) Secure Boot capable, c) TPM, and d) 64-bit compatible processor. I won't go into specifics about the tests I ran on my eight or so PCs running Windows 7 through Windows 10 other than to say that they all initially failed the test when using the flawed PC Health Check tool. It would appear that the clincher here is what Microsoft refers to as a "compatible 64-bit processor". For example, one may have a system where all the prerequisites are in place, except the processor does not meet the "a compatible 64-bit processor" requirement.

Microsoft has a list of what it currently deems as "compatible processors" for both Intel and AMD processors. Whether these two lists are regularly updated is anyone's guess. Going forward, unless Microsoft amends its list to include older processors which run Windows 10

perfectly well, it will turn a vast number of PCs, including their own Surface Pros, into zombies—machines frozen on Windows 10. Clearly, it's too early to know if Microsoft will release Windows 11 with relaxed requirements to support older or "non-compliant hardware", but it doesn't look that way. On the contrary, Microsoft seems to be digging in deeper on its position.

The compatibility issue will, without question, affect many *Maximum PC* readers, whose systems will likely not meet the Windows 11 requirements. My opinion on this issue is that Microsoft could release Windows 11 so that it gracefully degrades and runs on systems that are not fully "compatible" with Windows 11's new security features. Such an approach is possible because the current release of Windows 11 on the Microsoft Insider program will run on machines that do not strictly adhere to every requirement. So, as a reader and subscriber, I am hoping *Maximum PC* will be covering this contentious issue in the next issue and devote some serious time and copy to it.

–P. Sarro

EDITOR-IN-CHIEF, ZAK STOREY, RESPONDS: Thanks for the in-depth email, you're absolutely right, and it's something I want to address. In fact, there's a lengthy feature in this very issue concerning all things Windows 11, and I myself have written a neat little guide on how to get yourself a fresh install of Microsoft's latest OS.

I should perhaps preface that with the fact that this isn't actually a new OS as such. A lot of the core functionality, and code-base is actually the same as what we have on Windows 10. The majority of changes seem to be on a user interface level,

rather than from a new code perspective. There's new stuff here, particularly such as the Direct GPU to RAM interfacing, and support for various Linux and Android apps and such, and a whole swathe of new privacy features, for networking and security in particular, but on the whole, Windows 11, is just a glorified Windows 10. I've been using it for a while, and haven't encountered a single problem, with installing or testing anything apart from some slight UI bugs, which have all been patched out. Heck, even some of the debut features for Windows 11 have already been confirmed as being backported to Windows 10.

That gives me some hope, because what you say is true. The need for TPM 2.0, Secure Boot, and various other pre-requirements, although it does improve security, alienates a lot of those currently running Windows 10 on a myriad of devices. We also now have an end-of-life date for Windows 10 as well (October 14th, 2025), which is where Microsoft will no longer provide major updates.

That said, I agree they will inevitably have no choice but to renege on that requirement and allow older systems to upgrade to Windows 11. If for no other reason than national security for many countries across the globe. The number of systems running Windows 10 today in these government institutions must be phenomenal. I doubt they're all running the latest Ryzen and Intel chips, with TPM enabled as standard.

Asking all those governments and countries to replace every single PC in their infrastructure with new hardware, which would undoubtedly also require updates and patches to their own bespoke software, just because of a system requirement for Windows 11... well, that sounds like utter madness to me. ⏻

THE BUILDS

THIS MONTH'S STREET PRICES...

BUDGET

SO, THE TIME HAS FINALLY COME, we can upgrade our budget AMD's CPU. We have opted for the Ryzen 3 1200 CPU and seeing it in stock on Amazon.com was a joy, to say the least. The Ryzen 3 3100 would have been way better, but at around $345, it's more than double the price of the Ryzen 3 1200. At around $40 more than our previous CPU, we can now double our cores, but unfortunately, it won't double our threads just yet. Hopefully, in time, the Ryzen 3 3100 will slip into the budget build and make this a great full package.

Also in our AMD build, we hunted around and found some similarly-priced RAM sticks with a slight speed increase @ 3200MHz instead of @ 3000MHz. Crucial Ballistix 16GB (2x 8GB) DDR4-3200 CL16 Memory is a bit of a mouthful, but it will help us get that extra little bit of speed and future-proof it slightly. At $5 cheaper, we can't complain either. These changes help us improve our build whilst still not pushing the budget too far. When you are building your own custom rig, it's worth its weight in gold to look around. Yes, it can be time-consuming but, eventually, you'll get to know the market like the back of your hand.

Little changes here and there can soon equate to a decent saving, that you can spend more appropriately elsewhere. It's especially important seeing as our CPU has increased by $12 to $154. However, we have chosen the same RAM sticks as in the AMD build. For exactly the same price as our list issues RAM (Team T-Force Vulcan Z) we get the same configuration of 2x 8GB, but increase the speed from 3000MHz to 3200MHz. That makes sense, right?

The problem here for both builds is that GPUs are still awkward to get hold of. Prices have dropped slightly, but cards that you see floating around second-hand markets are still sky-high compared with the RRPs and aren't always the best. A lot of these may have been used for cryptocurrency mining, so it's like buying a two-year-old car with 150 miles on the clock, not exactly the best decision.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Corsair 4000D Airflow	$95	
PSU	500W EVGA BA 80+ Bronze	$40	
Mobo	MSI Pro B450 Pro-M2 MAX	$70	
CPU	AMD Ryzen 3 1200 **NEW**	$128	
GPU	AMD Radeon RX 6700 XT 12GB **NO STOCK**	$480	$800
RAM	16GB (2x8GB) Crucial Ballistix @ 3200MHz **NEW**	$75	
SSD	256GB Adata Falcon M.2 NVMe PCIe 3.0	$45	
HDD	1TB WD Blue 1TB 7200	$45	
OS	Windows 10 Home 64-bit OEM	$32	
Approximate Price:		**$982 or $1,302**	

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Corsair 4000D Airflow	$95	
PSU	500W ARESGAME AVG500 80+ Bronze	$35	
Mobo	Gigabyte B460M DS3H V2	$76	
CPU	Intel Core i3-10100	$154	
GPU	Nvidia GeForce RTX 3060 12GB **NO STOCK**	$330	$700
RAM	16GB (2x8GB) Crucial Ballistix @ 3200MHz **NEW**	$75	
SSD	256GB XPG Gammix S11 M.2 PCIe 3.0 **NEW**	$45	
HDD	1TB WD Blue 1TB 7200	$45	
OS	Windows 10 Home 64-bit OEM	$32	
Approximate Price:		**$887 or $1,257**	

MID-RANGE

AMD'S RYZEN 5 5600X still seems to be the best bet for our mid-range CPU. This 6 core processor with 12 threads is a great choice and thankfully, there's no price differential. We have opted for the Wi-Fi version of the same motherboard for an extra $15. However, a BIOS update may need to be performed for Vermeer CPUs to be compatible. We all know this can be a pain to do and you need an older CPU at hand that is supported by older BIOS versions.

Other than the change in motherboards, our mid-range AMD system hasn't seen any other changes. A second-hand Nvidia GeForce RTX 3070 is still pretty expensive and, if you fancy trying to get a new one, will probably be out of stock. Even finding a second-hand one for under $1000 is a mission and there is still the risk that it has been heavily used for cryptocurrency mining. G.Skill Ripjaws V Series RAM sticks are still a great pick at $85, and with a good speed, @3600MHz, they fit this build nicely. The cooler also drops down and saves us a whopping $3, Jolly Ranchers anyone?

A noticeable change is that the motherboard on our mid-range Intel build drops in price from nearly $190 down to $150. That's always a pleasant surprise and shows how unpredictable the market can be. For our CPU, we considered changing again to the i7-11700F counterpart. This is still more readily available and is now $70 cheaper. However, we stuck to our guns, as the chosen i7-11700K offers better performance and overclocking potential. The F variant also has no integrated graphics, but if you aren't put off by that, it's still a great CPU, especially now it's cheaper.

This time we can pick up the Enermax Liqmax III 240 for $60—we are always happy to see a price drop. We also opted for the GA version of the last issue's PSU as it's virtually the same unit but with more connectors. Last time, we added Corsair's Force MP600 M.2 SSD, which was a great addition. This PCIe 4.0 SSD has great speeds and still can be found for $100, which is a definite pick-up.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Lian Li PC-O11-Dynamic	$150	
PSU	650W Fractal Design Ion Gold	$105	
Mobo	Asus AM4 TUF Gaming X570 -Plus WiFi **NEW**	$195	
CPU	AMD Ryzen 5 5600X	$290	
Cooler	240mm Cooler Master Masterliquid ML240L RGB	$67	
GPU	Nvidia GeForce RTX 3070 8GB **NO STOCK**	$500	$1,065
RAM	16GB (2x 8GB) G.Skill Ripjaws V Series @ 3,600MHz	$85	
SSD	500GB PNY XLR8 CS3040 M.2 PCIe 4.0	$100	
HDD	1TB WD Blue 1TB 7200	$45	
OS	Windows 10 Home 64-bit OEM	$32	

Approximate Price:	**$1,569 or $2,134**

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Lian Li PC-O11-Dynamic	$150	
PSU	650W EVGA SuperNOVA 650 GA **NEW**	$80	
Mobo	ASRock Z590M PRO4	$150	
CPU	Intel Core i7-11700K	$400	
Cooler	Enermax Liqmax III 240 RGB	$70	
GPU	AMD Radeon RX 6800 XT 16GB **NO STOCK**	$650	$1,300
RAM	16GB (2x 8GB) Silicon Power XPOWER Turbine @ 3,600MHz	$85	
SSD	500GB Corsair Force MP600 M.2 PCIe 4.0	$100	
HDD	1TB WD Blue 1TB 7200	$45	
OS	Windows 10 Home 64-bit OEM	$32	

Approximate Price:	**$1,762 or $2,412**

TURBO

YES, HERE WE ARE AGAIN talking about GPUs in a negative way. It seems such a shame but this elephant in the room doesn't seem to be going away any time soon. Your best bet to finding a new graphics card is by keeping yourself right up to date by tracking restocks via social media updates and camping out on reseller websites.

Once you are in these groups, you will hopefully get a heads up and at least have a chance to buy from new. We can't recommend paying for a second-hand market GPU, as the prices are still relatively high and you can never guarantee the quality you are getting. So in this case, both GPUs are still out of stock.

This time around, we have swapped out the Patriot Viper Steel DDR4 RAM on the AMD build for some new sticks. G.Skills Ripjaws V sticks offer faster speeds @4000MHz and are still the same 2x 16GB configuration for just $8 more. The central processing chip stays the same here. When looking around for a similar price, the Ryzen 7 5800X CPU is still our best bet. With 8 cores, 16 threads, and a 3.8GHz base clock, it's a great gaming CPU.

The 980 Pro Samsung SSD has seen a slight price drop, making it an even better choice. At the time of writing, a Newegg promo gives you an extra 15 percent off, so at $170, there's no better time to pick up this SSD. With some cuts in the budget, we upgraded our AIO to Corsair's iCUE H150i ELITE CAPELLIX 75 cooler. This costs around $25 more but is a great all-around 360mm AIO.

On the Intel rig, another sale on Newegg brings our newly added mobo down by $25. The beautiful i9-11900K remains on the turbo build and also sees a discount. We don't think this will be swapped out any time soon and is a no-brainer for the Intel system. The same RAM was also added to this build, as seen in the AMD turbomachine. As this RAM is cheaper than the previous sticks, it makes sense to pick the same 2x 16GB configuration just with increased speeds. Other than that, the Intel build has no changes, just a decrease in prices across the board, apart from the GPU unfortunately. On the whole, though, it's a positive direction for the market.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$150	
PSU	750W NZXT C750 80+ Gold	$105	
Mobo	MSI MPG X570 Gaming Edge Wi-Fi	$211	
CPU	AMD Ryzen 7 5800X	$397	
Cooler	Corsair iCUE H150i ELITE CAPELLIX AIO 360mm **NEW**	$170	
GPU	Nvidia GeForce RTX 3080 10GB **NO STOCK**	$700	$1,750
RAM	32GB (2 x 16GB) G.Skill Ripjaws V @ 4000 **NEW**	$163	
SSD	1TB Samsung 980 Pro M.2 PCIe 4.0	$170	
HDD	6TB Western Digital Blue	$150	
OS	Windows 10 Home 64-bit OEM	$32	

Approximate Price: **$2,248 or $3,348**

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$150	
PSU	750W EVGA SuperNOVA 750 G5	$100	
Mobo	ASRock Z590 Steel Legend Wi-Fi	$195	
CPU	Intel Core i9-11900K	$530	
Cooler	MSI MAG CoreLiquid 360R AIO	$120	
GPU	AMD Radeon RX 6900 XT 16GB **NO STOCK**	$999	$1,475
RAM	32GB (2 x 16GB) G.Skill Ripjaws V @ 4000 **NEW**	$163	
SSD	1TB Gigabyte Aorus Gen4 7000s M.2 PCIe 4.0	$200	
HDD	6TB Western Digital Blue	$150	
OS	Windows 10 Home 64-bit OEM	$32	

Approximate Price: **$2,639 or $3,115**

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

LOVE
PAPER ®

We build the world's most advanced PCs.

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

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM
Digital Storm PCs featuring Intel® Core™ i7 processors.

(intel) CORE i7 inside