



### GTX 1660 SUPER

Is this the new affordable GPU of choice? **PG. 80**



### VIRTUALBOX TRICKS

Speed up your VMs and add multiple displays **PG. 58**



### MASTER AUDIOBOOKS

Rip, convert, & stream to all your devices **PG. 46**

# MAXIMUM PC

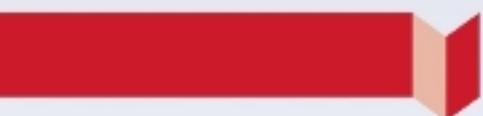
MINIMUM B5 • HOLIDAY 2019 • [www.maximumpc.com](http://www.maximumpc.com)

**BUILD IT**

# 4K GAMING MONSTER

- ✓ Create a custom cooled masterpiece
- ✓ Dual SLI GeForce RTX 2080 Supers
- ✓ \$6K of powerful PC performance

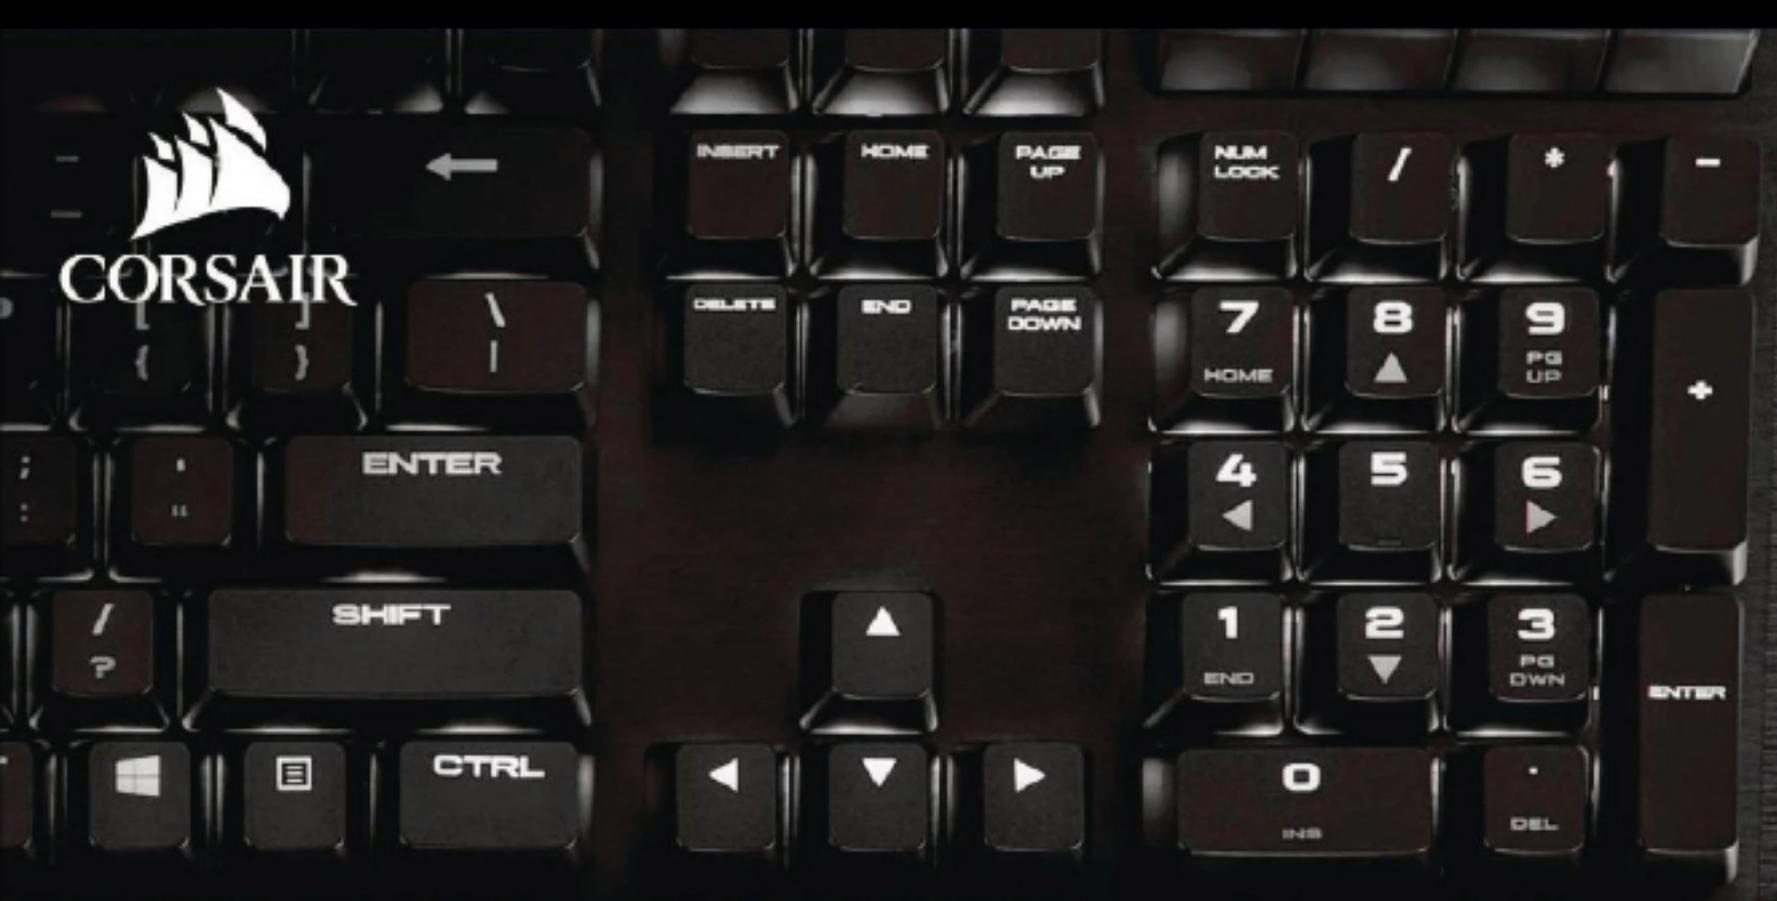
**NO.1**  
FOR PC  
HARDWARE



## INTEL 10TH-GEN CPUs EXPLAINED

Discover the real gems of the latest chip family **PG. 38**

FUTURE



INCRECIBLE SOUND

IMPECCABLE QUALITY

# VIRTUOSO RGB WIRELESS

## High-Fidelity Gaming Headset

The CORSAIR VIRTUOSO RGB Wireless delivers a high-fidelity audio experience for the most discerning players, combining uncompromising sound quality with all-day, all-night comfort from its premium memory foam earpads and headband.

Connect to virtually any device with ultra-low latency SLIPSTREAM CORSAIR WIRELESS TECHNOLOGY, 3.5mm wired, or 24bit/96KHz USB wired. Make yourself heard with a broadcast-grade, omnidirectional detachable microphone with exceptional dynamic range.

find out more at [corsair.com](http://corsair.com)

# GIFT YOURSELF

## TRUE GAMING PERFORMANCE

You Deserve the Best **4K Ultra Settings** in Every Game

Personalize your laptop with **HD UV Printing**



EVO16-S

NEURON



Powered by Intel® Core™ i9 Processors

**FREE** Lifetime 24/7 Service & Support  
Online Financing Available

Intel, the Intel logo, Intel Inside, Intel Core are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

# STUFF YOUR NEW PC NOT YOUR STOCKINGS

Fully Customizable with  
Patented Variable Mounting



MILLENNIUM

[ORIGINPC.COM](https://originpc.com)

@originpc [f](#) [i](#) [t](#) [v](#) [t](#)

**ORIGIN**  
HIGH-PERFORMANCE PCS

© 2019 ORIGIN PC Corporation. All rights reserved. ORIGIN "O" Symbols are trademarks of ORIGIN PC Corporation.  
All other trademarks are property of their respective owners.

**SUBSCRIBE TODAY!**  
see PG. 44



# inside

## HOLIDAY 2019

### QUICKSTART

- 10 THE NEWS**  
Streaming wars; Flash nearly dead; Twitter goes apolitical; Win 7 due date; EA back on Steam; more....
- 16 THE LIST**  
The eight best G-Sync monitors available right now.
- 18 HEAD TO HEAD**  
PCIe Gen3 M.2, PCIe Gen4 M.2, or SATA SSD—which wears the storage crown?



Dream screen: the Asus ROG Swift PG27UQ.

### R&D

- 56 AUTOPSY**  
Does the Fairphone 3 give us a fair stab at repairs and replacements?
- 58 HOW TO**  
VirtualBox tips; monitor tweaks for gamers; Luminar 4; push live system stats to the cloud; edit video for free; enhance soundtracks.

### LETTERS

- 20 DOCTOR**
- 94 COMMENTS**

**22 LIQUID-COOLED MONSTER**  
How we put together our \$6,000 Ryzen 4K SLI gaming beast.

**38 10 IS THE MAGIC NUMBER**  
Intel's 10th-gen Core CPUs are here, but the power of 10 is proving a little problematical....

**46 TAKE CONTROL OF YOUR AUDIOBOOKS**  
Rip, convert, and stream your collection to any device.

### IN THE LAB



**76 FALCON NORTHWEST TALON**

**80 NVIDIA GEFORCE GTX 1660 SUPER**



**86 TERRAMASTER F5-422**

**90 BORDERLANDS 3**



© ASUS, 2K GAMES





# MAXIMUMPC

## EDITORIAL

**Executive Editors:** Alan Dexter  
**Senior Editors:** Jarrad Walton  
**Hardware Lead:** Ito Moore  
**Hardware Staff Writer:** Joanne Nalua  
**Staff Writers:** Christian Guyton  
**Contributing Writers:** Alex Campbell, Alex Cox, Ian Ewenden, Phil Iannak, John Knight, Jeremy Lord, Chris Lloyd, Carrie Marshall, Alice Newcome-Bell, Nick Peters, Les Pounder, Zak Storey  
**Copy Editors:** Katharina Davies  
**Editor Emeritus:** Andrew Sanchez

## ART

**Art Editor:** Fraser McDermott  
**Photography:** Phil Barker, Billy Curtis, Neil Gordon  
**Cover Photo Credits:** Future plc, Nikita Oracki

## BUSINESS

**US Marketing & Strategic Partnerships:** Stacy Gaines, stacy.gaines@futurenet.com  
**US Chief Revenue Officer:** Luke Edson, luke.edson@futurenet.com  
**East Coast Account Director:** Brandie Rushing, brandie.rushing@futurenet.com  
**East Coast Account Director:** Michael Plump, michael.plump@futurenet.com  
**East Coast Account Director:** Victoria Sanders, victoria.sanders@futurenet.com  
**East Coast Account Director:** Melissa Planty, melissa.planty@futurenet.com  
**East Coast Account Director:** Elizabeth Reichman, elizabeth.reichman@futurenet.com  
**West Coast Account Director:** Austin Park, austin.park@futurenet.com  
**West Coast Account Director:** Jack McAuliffe, jack.mcauliffe@futurenet.com  
**Director, Client Services:** Tracy Lam, tracy.lam@futurenet.com

## PRODUCTION

**Head of Production:** Mark Constance  
**Production Manager:** Vivienne Calvert  
**Project Manager:** Clare Scott  
**Production Assistant:** Emily Wood

## FUTURE US, INC.

11 West 42nd Street, 15th Floor, New York, NY 10036, USA  
[www.futurenet.com](http://www.futurenet.com)

## SUBSCRIBER CUSTOMER SERVICE

Maximum PC Customer Care,  
 Future Publishing, PO Box 9952, Harlan, IA 51593-1952  
 Website: <http://www.futurenet.com>  
 Tel: 866-779-2877  
 Email: [contact@myfavortemagazines.com](mailto:contact@myfavortemagazines.com)

## BACK ISSUES

Website: <http://www.futurenet.com>  
 Tel: 465 341 808 2852

Next Issue On Sale January 9, 2020

© 2019 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.



Future plc is a public company quoted on the London Stock Exchange (symbol: FTSE)  
[www.futureplc.com](http://www.futureplc.com)

Chief executive: **Edith Cowley-Thorne**  
 Non-executive chairman: **Michael Hastings**  
 Chief finance officer: **Penny Liddle-Grind**

Tel: +44 (0)1226 443 244



Alan  
Dexter

## SOME DREAMS SHOULD BE ATTAINABLE

**LET'S CUT TO THE CHASE:** This month's cover build is essentially a Dream Machine. It's beautiful, packed with some of the latest hardware, and is a total monster in terms of performance. Having said that, when it comes to component selection, we've kept a close eye on value for money more than our Dream Machine traditionally does, so while, at \$6,000, it isn't exactly cheap, it doesn't cost the sort of money that would normally net you a car or a new room for your house.

So, no, officially the system we've built this issue isn't actually a Dream Machine, but in terms of performance, it's very close, and it's an incredible PC all the same. If you've ever wanted to produce the kind of machine that you could show off to your friends and family, this is it. Boasting custom dual water-cooling loops, stunning attention to detail, and a final design that wouldn't look out of place at most computer fairs, this really is a machine to behold and show off. It isn't what anyone would really call easy to put together, but with enough patience, money, and our help, you can achieve your own monster PC.

In case you're wondering, this system uses an AMD CPU, the Ryzen 9 3900X, as we see it as the current pinnacle of consumer desktop processing (at least until the Ryzen 9 3950X drops—at the time of going to press, AMD has just revealed the final details of this chip, along with info about the next-gen Threadrippers, which we'll cover in more detail next issue). But what of Intel? This month, we take a look at its 10th-generation processors. They're predominantly laptop focused, but there are some offerings for desktop as well—although, as ever,

it isn't straightforward, and with multiple architectures and production processes in play, it pays to know what you're paying for before you pull the trigger. Turn to page 38 to find out more.

We talk a lot about managing music and videos here in *Maximum PC*, but one area that often gets overlooked is audiobooks. This issue, we correct that, with a complete guide to taking control of your audiobooks: pulling your books from various services and saving them locally, and organizing them in a way that makes sense. We also look at other sources for audiobooks, including free options that you may not know about.

In our tutorial section, you'll find plenty of helpful advice for doing more with the hardware you already have, kicking off with VirtualBox hints and tips. The rise of Docker may mean that its one-stop containers can be the way to go for specific tasks, but taking complete control of the system for more general use still works in the favor of traditional virtualization, and being a master of that will help you get more things done. This is followed by a tutorial on how to get the most from high-refresh screens, and why just leaving everything at the whim of your graphics driver means that you won't be getting the best possible experience.

I hope you enjoy the issue!

Alan Dexter is *Maximum PC's* executive editor and a punisher of hardware. He's been a tech journalist for over 20 years, and has no problem upsetting the PC industry as a whole.

submit your questions to: [comments@maximumpc.com](mailto:comments@maximumpc.com)

# THE NEWS

## Streaming Wars

The fight is on for your video streaming subscriptions

**NEXT YEAR** will see the start of a major fight between the tech and media giants over the video streaming market, and huge amounts of money are about to be thrown at your screen to earn your subscriptions. Network TV is on the decline—a generation has appeared that doesn't consume media that way anymore. Why wait when you can binge? Why use cable when you've got the Internet? Is live TV dead? No, there's always room for sport and news, and it's also making the jump to the Internet. Some streaming services include live television channels, and there are also exclusively live TV services, aimed at cable-cutters. The good news is that the fierce competition means there are going to be some very good deals.

Netflix has lead the push into streaming, and become a Fortune 500 company in the process. It has about 140 million subscribers. Amazon Prime has, well, no one outside Amazon is sure how many subscribers it has, as its video streaming is bundled with other services, and Amazon doesn't share its numbers,

but there are over 100 million Amazon Prime members. These are the two big players, and so far Netflix has lead the field, with lots of decent original content, including some that wins Oscars.

The Motion Picture Association estimates there are 613 million subscribers worldwide, in a market worth \$22 billion, and estimated to grow to over \$30 billion by 2022. This already competitive market is about to get a good shaking due to the addition of two more giants this November: Disney and Apple.

Apple TV+ will cost \$4.99 a month, undercutting the competition. Apple will also be throwing in a free year's subscription if you buy a new piece of Apple gear, so the initial subscriber base is going to grow very rapidly. Where Apple will struggle is content; it has signed some big names, but it doesn't own the big back catalogs that many of its rivals have. However, Apple is rich, and this is the kind of problem that can be cured with money.

Disney bought 20th Century Fox this spring, which also brought Hulu into its stable. This will be run alongside its



**Fight!** All these, and many more, want your money. There is going to be an awful lot of video to watch next year.

new Disney+ service carrying more adult material, while the Disney brand keeps the family-oriented material. It has an impressive store of content, including Disney cartoons, Pixar animations, as well as the Marvel and Star Wars franchises. Disney+ costs \$6.99, while Hulu will be bundled with ESPN+ at \$11.99 without ads, or \$5.99 with ads.

In spring, HBO will expand its operations with HBO Max. This \$14.99 service has access to a massive library from WarnerMedia. NBCUniversal is also joining the fray with Peacock, due in April at an as yet undisclosed price. It has access to the combined back catalog of NBC television and Universal Pictures.

Every big media company wants a slice of the Netflix market, but Netflix doesn't make any profit. The tactics used by all the players are similar: Keep prices competitive, get the numbers

up, and invest in content. Top-tier content is hideously expensive; Amazon is to put a reported \$500 million into its *Lord of the Rings* series. Profit is something you make much later—HBO Max isn't expected to make any until 2025.

There's over 100 streaming services in the US, offering everything from live TV, to specialist on-demand content, to media giants with everything. A lot of choice—and a lot of subscriptions. People seem happy to pay for one or two subscriptions, but how about three, or four?

There has already been one victim. Sony's PlayStation Vue is to close in January. Sony said it has "decided to remain focused on our core gaming business." It was a tough sell at \$45 a month. This will get ugly, and there will be blood. However, we'll get a lot of top-notch content, at reasonable prices, if you can decide which service to use. **-CL**



Why wait when you can binge? Why use cable when you've got the Internet?

# GET THE MOST FROM WINDOWS 10 WITH OUR COMPLETE GUIDE

All that you need to make the most of your Windows PC or laptop,  
including the best tips, tricks and tutorials!



Ordering is easy. Go online at:

[www.myfavouritemagazines.com](http://www.myfavouritemagazines.com)

Or get it from selected supermarkets, newsagents, or bookstores



## TWITTER GOES APOLITICAL

Media giant to ban all political advertisements

**TWITTER IS TO BAN** all political and interest advocacy ads. The news was broken by Twitter's CEO, Jack Dorsey, in a Twitter thread. He said that "forcing" highly optimized and targeted messaging "brings significant risks to politics." He says that this is a power that "should be earned, not bought," and "this isn't about free expression." The only ads Twitter will allow are ones for voter registration. Twitter's decision will please some, although it'll be difficult to enforce completely; "political" can be difficult to define.

What of Facebook? The company has been under heavy fire for not doing enough to remove disinformation. Mr. Zuckerberg cited the First Amendment to defend the company's policy, free speech being held as paramount. Facebook is officially neutral, which has allowed highly spurious content to go unchecked. Zuckerberg stated recently that "In a democracy, I don't think it's right for private companies to censor politicians or the news."

Over 250 Facebook employees signed a letter to Zuckerberg saying the policy was a "threat to what Facebook stands for." The letter proposes that political ads be clearly marked and limited in their targeting, and imposing a spending cap.

Dorsey and Zuckerberg don't get on. The two founded companies with quite different approaches to user anonymity, and how to use their platforms. Neither is beyond taking a potshot at the other's stance. The use of political messages in social media is contentious. So many, especially the young, use such platforms to gather news, and the ability to use AI learning to micro-target people with tailored ads is a powerful tool. The lack of regulation leaves both companies' founders free to decide what is acceptable. We can expect these two to continue to snipe at each other for a while yet. **-CL**

## FLASH NEARLY DEAD FROM HERO TO ZERO



**GOOGLE HAS ANNOUNCED** that it is no longer index Flash files for its search engine, making Flash content invisible to most people. A little premature perhaps, but Flash is dying quickly now. Adobe is due to stop distributing and updating it before the end of next year. The latest version of Chrome has it disabled by default, and the next version won't run it at all. Despite this, of the top 1,000 websites, over 8 percent sport Flash content.

Why is it dying? Flash is resource-heavy, and uses elderly video decoding, so it drains phone batteries quickly. It is also propriety, i.e. not free, unlike the better open standards that followed. It's vulnerable to hacking, too, and Adobe was sluggish in fixing those vulnerabilities. In its day, Flash was everywhere, adding animation to the static websites of the '90s, and powering a generation of online games. It also enabled Internet upstart YouTube to distribute videos. Flash will live on offline, as an animation tool, but soon you won't be able to find, let alone play, Flash files online easily. **-CL**

## WIN 7 END DATE

The nagging starts in earnest



**THE DAYS OF WIN 7** are numbered: Extended support is due to end on January 14, 2020. What happens after the end of support? Not much—it will continue to work, but there will be no official support. More importantly, there will be no security updates. Probably. If something really nasty happens, Microsoft can issue an "emergency" patch (it has done this before).

If you're running Office 365, you'll still get security updates until January 2023, but you won't get any new features or improvements. Other versions of Office will get bug fixes for their planned lifespan, which runs from October 2020 for Office 2010 to October 2025 for Office 2016. Internet Explorer 11 is less lucky; it'll stop being updated along with Win 7. If you really must run a Windows 7 PC, Microsoft will sell you an Extended Security Update (commercial customers only). Current estimates have Win 7's market share at around 35 percent. That's a lot of PCs (and a lot of upgrades to sell). **-CL**

## Tech Triumphs and Tragedies

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

### TRIUMPHS

#### NEW SONY PSVR?

Some interesting patents show Sony is working on a new wireless VR headset for the PS5.

#### BIGGEST BATTERY EVER

New York is to build a 316MW lithium-ion storage battery, enough to power 250,000 homes for eight hours.

#### \$600M IN THREE DAYS

*Call of Duty: Modern Warfare* had a record-breaking first weekend and is the top-selling premium game of the year.

### TRAGEDIES

#### BRICKED HOMEPODS

An OS update to Apple's HomePod caused some to crash irrevocably; the update has since been pulled.

#### GOT YOUR NUMBER

Google and Facebook have been extracting, indexing, and tagging license plate numbers from photos without asking.

#### REMOTE CONTROL CAR

Months after returning a rented Ford Expedition car, a man can still control it from his phone.

## Electronic Arts Back on Steam

**EA HAS ANNOUNCED** that it plans to migrate its games back to Steam, starting with *Star Wars Jedi: Fallen Order*. Other titles will make the jump over the next few months. You still need an EA account, and Origin is still part of the setup process, for now at least. EA is also bringing its EA Access subscription service to Steam.

Steam wasn't very popular when it appeared back in 2003. Initially used to update Valve games, that all changed when it became mandatory if you wanted to play *Half-Life 2*. At which point, the servers crashed, and there was a lot of online moaning. By 2007, it had brought most of the biggest players in the business onboard. These days, it dominates game distribution, and has earned its creators an estimated \$4 billion.

However, Electronic Arts was a noticeable absentee. It launched its own distribution service in 2011. Intended as a direct Steam rival, Origin never really gained enough support outside EA. This relegated it to an extra requirement if you wanted to play EA's top titles, which is tiresome when the rest of your library is on Steam.

Why the thaw? EA vice president Mike Blank summed it up: "We want to be where the players are." Steam is too big to ignore, but it's not immune to competition. Epic Games has its own service, Epic, launched last year, built on the back of the success of *Fortnite*, which earned the company an immense pile of cash and millions of players.

Steam takes a 30 percent cut of game revenue, Epic will take 12 percent. Given that profit margins on a game are a fraction of the price, that is a big incentive for publishers to jump ship. Steam has responded by reducing rates a little if you have sales of over \$10 million, but its fat margins have attracted a well-funded rival. **-CL**



*Star Wars Jedi: Fallen Order*, the first big Electronic Arts title on Steam for eight years.

## Google Buys Fitbit

Fitbit has not had a good year. Competition from cheaper rivals, production troubles, and a failure to break into the lucrative smart watch sector has led to financial trouble and a plunging stock price. Google has stepped in and bought the company for \$2.1 billion, 70 percent more than the stock price. Google has its own fitness tracker, but previously relied on others for the hardware. Now it has its own hardware, and 28 million active users. Meanwhile, Fitbit gets all the software resources it needs. The obvious question is: Who now owns your data? Fitbit issued a statement that "health and wellness data will not be used for Google ads." If you have a Fitbit, you will be given the option to delete your data. That's some comfort at least. **-CL**

## Atom Gets New Microarchitecture

Intel's ultra-low voltage x86 CPU, the Atom, is getting a new core, code-named Tremont. The 10nm chip has "Core class" branch prediction with out of order fetch. The front end boasts a pair of decode clusters, each with decode units, to enable parallel out-of-order decode, and there is up to 4.5MB of L2 cache. Intel reports that performance of early samples is 30 percent or better than the previous Goldmont Plus design. The Tremont core is due to appear first in the hybrid notebook Lakefield chips late this year. This stacked design puts a Sunny Cove core and four Tremont Atom cores, along with Gen11 graphics, the I/O, and cache, all in one neat layered package. **-CL**



## xCLOUD STREAMING GOES BETA

Rival to Google's Stadia is on its way

**MICROSOFT HAS** started signing up players for a public beta test of its xCloud service. Like Google's Stadia, xCloud will render games on racks of servers (basically lots of Xbox machines), and stream the screen image back to your device. You can play demanding top-tier games on fairly humble devices, as long as your Internet connection can handle it. That's the theory anyway. It's not a new idea, the main problem has always been latency: Can you get the right level of responsiveness while running a game remotely? Reactions are pin-sharp on modern consoles, and players will expect the same on xCloud.

To participate in the beta test you need a device running Android 6.0 and Bluetooth 4.0 or better, and a Bluetooth controller, typically an Xbox controller, although not all are compatible. You also need a connection of 10Mb/s or better. There are three initial games: *Sea of Thieves*, *Gears 5*, and *Halo 5: Guardians*. A decent test of the system's capabilities.

Microsoft has another streaming service in beta, Xbox Games Streaming, which does a similar trick, only your game is streamed from your Xbox to an Android device. Otherwise the requirements are similar, but it needs 4.75Mb/s upstream bandwidth, and a latency of 125ms or under. You can now play Xbox games remotely on your smartphone, although original Xbox and Xbox 360 games are not currently supported.

xCloud goes head to head with Google's Stadia, which should have dropped by the time you read this—the official launch date is November 19. The list of games announced so far runs to 43, and includes exclusives, along with a few currently unreleased games (including *Baldur's Gate 3*). This should be good. **-CL**

**“ Curiosity about  
life in all its  
aspects, I think, is  
still the secret of  
great creative people.**

---

**LEO BURNETT**

---

The number one destination for  
**graphic design news, views, and how-tos**

---



Jarred Walton

## TECH TALK



# Do You Really Need to Upgrade Your CPU?

Just when things seemed to be settling down, a whole new wave of processors is descending upon us. AMD's third-generation Ryzen CPUs are proving to be potent competitors to their Intel counterparts, and while 10nm mobile processors from Intel are shipping, the battle for the desktop is beginning to look grim.

But Intel isn't ready to concede just yet, which is why it launched the Core i9-9900KS—a CPU that can legitimately clock at 5.0GHz on all cores without breaking a sweat.

That's the least of the news, however. The i9-9900KS is just a pre-emptive strike against AMD's continuing onslaught. During the next month, I'll be looking at no fewer than four additional extreme CPUs, all vying for your cash. First is AMD's Ryzen 9 3950X, the 16-core/32-thread processor that's wrapping up the mainstream socket AM4 family. In an interesting twist, this chip won't include a cooler, and AMD recommends a liquid-cooling solution, so it's a lot like Intel's K- and X-series parts. Next will be Intel's Cascade Lake-X, the third round of 14nm++ Skylake-X parts, topping out yet again with an 18-core/36-thread part, the Core i9-10980XE. There will be other chips fleshing out the family, but the big news is that Intel has slashed prices in half relative to its ninth-gen parts.

That's because AMD is forcing Intel's hand, with a third generation Threadripper family on a new TRX40 platform. It's like socket sTR4, except with additional PCIe lanes—plus the PCIe lanes are all Gen4. The link to the chipset is also an x8 Gen4 connection instead of x4 Gen3, so four times the bandwidth. The cherry on top is that AMD will again offer 32-core/64-thread and 24-core/48-thread CPUs, except this time I expect performance will

be much better than the second-generation Threadripper parts, thanks to the refinements AMD has made in the Zen 2 architecture.

It's an exciting time for CPUs, in other words, but it does bring up an interesting question: Who should be looking to upgrade to the latest and greatest parts? There's no single correct answer, of course.

Some people want the best, and they'll happily throw down for the fastest CPU on tap. Others only want to upgrade when their current rig becomes too slow—which might mean a chip that's only one or two years old, or it could be a Core i5-2500K CPU from 2011, or anything in between. I can't tell you whether you really need to upgrade, but I can tell you what to expect in terms of performance improvements.

The fastest mainstream CPU from two years ago is the Core i7-8700K, which could usually hit 5.0GHz with overclocking. The new i9-9900KS runs 5.0GHz "stock" (and can overclock to 5.2GHz with a good cooler), and it's up to 55 percent faster in some workloads. The catch is that in other workloads, such as games and everyday office tasks, the difference is far less pronounced. Even with the fastest current graphics card, an RTX 2080 Ti, you'll only see about a 5 percent

difference in game frame rate on average—and that's at 1080p.

Compared to a slightly older Core i7-6700K, with half the cores and threads, never mind the lower clock speeds, performance in some tasks can be more than twice as fast. Gaming performance, though? It's still only about 15 percent faster on average. Even a Core i7-4770K won't really put the hurt on gaming performance. The i9-9900KS is nearly 50 percent faster at 1080p ultra, but only 10–15 percent faster at 1440p ultra—which is what you'd likely be using with such an extreme GPU.

But for the subset of users who always want more CPU performance, hang on to your wallets. Intel may have chopped prices on its HEDT lineup in half, but AMD is filling the gap. The 24-core Threadripper 3960X will sell for \$1,399, while the big behemoth 3970X takes over Intel's former \$1,999 stomping grounds. I'm still looking forward to seeing it chew through multithreaded workloads such as Cinebench, even if it will probably trail the 9900K in gaming performance. It's a wonderful time to be a hardware enthusiast.

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



Intel may have chopped prices on its HEDT lineup in half, but AMD is filling the gap.

# THE LIST

## THE BEST G-SYNC MONITORS RIGHT NOW

8

### ASUS TUF GAMING VG27AQ

Whether you buy into the TUF series's stripped-back aesthetics is a matter of taste, but this is a solid G-Sync IPS monitor.



4

### ACER PREDATOR XB321HK

This G-Sync monitor is reasonably priced for a 4K screen, thanks to its lack of superfluous HDR features.



7

### DELL S2719DG

A cheaper option that might not look as flashy as others, but gets the job done well with anti-glare and 1ms response time.



3

### ASUS ROG SWIFT PG27UQ 4K,

144Hz, and a ton of ports make this one of the best G-Sync monitors, but the price hurts at close to \$2,000.



6

### AOC AGON AG352UCG6

Excellent G-Sync performance on a curved display, plus fantastic color contrast with a blue-light mode for night-time gaming sessions, too.



2

### ACER PREDATOR XB273K

Yes, the screen shroud is weird, but this 4K display does out stunning graphics without too harsh a price tag.



5

**ALIENWARE AW3418DW** It's expensive, but this ultrawide beauty packs a 100Hz refresh rate—overclockable to 120Hz—with a sleek design to boot.



1

### ASUS ROG SWIFT PG279Q

Overclockable to 165Hz, this could be the best 1440p monitor available. It's reasonably priced, super-responsive, and will last ages.





Alex Campbell

## OPEN SOURCE

# Why We Respect Ubuntu

**WHETHER YOU LOVE** Canonical's Debian-based Linux offering or loathe it, Ubuntu has an important role to play in the Linux ecosystem as one of the most popular OSes for desktop and server applications. With Ubuntu 19.10 hitting the net back in October, it might be good to think a bit about what Ubuntu's release cycle provides for the Linux ecosystem.

October means a lot of things in the United States. It marks the month of Halloween and the beginning of the holiday season. But for Linux users, it also marks another annual tradition: the release of a new version of Ubuntu.

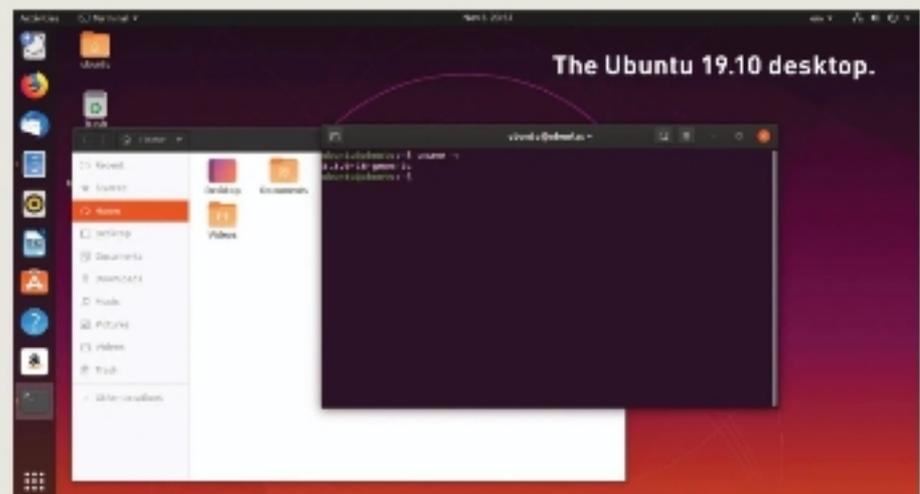
Just as Intel has long followed the tick-tock release cycle for its CPU architectures, Ubuntu follows a regular rhythm for its release. But unlike a simple tick-tock, Canonical releases Ubuntu in what feels like a 4/4 time signature in music, with each release acting as a beat, or quarter note.

For those not familiar with Ubuntu versions, it's fairly easy to know when the releases were pushed. Versions are numbered with a two-digit year, followed by a two-digit month, and optionally the letters "LTS." For example, the latest release is numbered 19.10, signifying October of 2019.

Every even year, Canonical pushes out a long-term support (LTS) version of Ubuntu in April. This LTS release can be considered the downbeat, or beginning of the measure. The last downbeat (1/4) was in 2018, with 18.04 LTS. The immediate offbeat (2/4) that year was 18.10, which was followed by the next downbeat (3/4) with 19.04. The cycle finished with the last beat (4/4) in October, with 19.10.

For most users, I recommend sticking with LTS releases, which are supported by Canonical for five years. However, going with one of the other releases can be fine, though there may be more bugs. Most non-LTS releases—especially the odd October releases—serve as testing grounds for features that will appear in the next LTS release.

There's a reason for this kind of schedule. Sure, Ubuntu may not officially get to use new kernels as soon as they're available. And yes, sometimes versions for certain packages may be a version or more behind what you'd find in a rolling-release distribution, such as Manjaro or OpenSUSE Tumbleweed. But what Ubuntu's LTS releases provide is a system for developers to target,



considering many distros (such as Mint) are based on Ubuntu, the LTS target makes even more sense.

All well and good, but one thing stands out about non-LTS Ubuntu versions: They only receive nine months of support and updates. For people running businesses or servers, that's not acceptable, even if it's fine for your average user who's happy with upgrading the system every six months.

Although Ubuntu's LTS versions are a good target, Ubuntu doesn't always play nice with Debian, the distro Ubuntu is based on. While Ubuntu can generally use packages designed for Debian (thus the .deb package format), if you try to use packages designed for Ubuntu on Debian, you're playing with fire. This is a sore spot for some in the Linux community, and one among many gripes that some people have with Canonical. But one can't deny

Ubuntu's role in creating an easy-to-install-and-use desktop distro that brings new users to Linux.

Should everyone use Ubuntu? Of course not. But it can be a good metric. If you want something more stable, try Debian or Red Hat Enterprise Linux. Fedora follows a six-month release cycle, too, but deprecates each version after 13 months. If you want the latest security updates, kernels, and features, a rolling-release distro based on Arch or OpenSUSE Tumbleweed will serve you better (if you're OK with the risk of the odd update breaking something).

If you're building a workstation or home server, or are preparing someone's first Linux system, an Ubuntu LTS is my first choice.

Alex Campbell is a Linux geek who enjoys learning about computer security.

## HEAD TO

BY CHRISTIAN GUYTON

# PCIe Gen3 M.2 vs. PCIe Gen4 M.2 vs. SATA SSD

The fourth generation of peripheral component interconnect express—PCIe, for those of us without time to waste—has arrived, and more Gen4 M.2 drives are appearing. But is it truly superior to PCIe Gen3? The fourth generation is still in its fledgling stage, while the range of Gen3 M.2 drives is expansive and varied. Then there's their predecessor: good old SATA SSD. Which deserves the storage crown? (No, that's not a real thing.) Let's find out....

## ROUND 1

### Value

Uh-oh! PCIe Gen4 stumbles right out of the gate. While the new M.2 SSDs boast incredible specs, those specs come with a less-than-excellent price tag. Even the cheapest Gen4 SSD is going to set you back over \$100, with larger sizes, such as 2TB drives, running over \$400. The performance gain over Gen3 simply isn't significant enough to justify the price shift, unfortunately. Meanwhile, Gen3 M.2 prices have dropped since the release of the new generation, meaning it's now a great time to make a purchase. A terabyte of high-speed Gen3 M.2 storage is currently available for scarcely more than \$100—a far cry from the expensive HDDs of yesteryear.

You might expect SATA SSDs to be the cheapest, and you'd be absolutely right—but not as right as you may think. While low-capacity SATA drives (120GB and below) are still an absolute bargain, the price difference between 1TB SATA and Gen3 M.2 drives is easily measurable in double digits, not triple. Given the huge performance differences between the SATA and M.2 formats, that's not a big gap. While small SATA SSDs are useful as boot drives for budget systems, overall it's difficult to recommend building a new PC without including an M.2.

**Winner: PCIe Gen3 M.2**

## ROUND 2

### Performance

A surge back into pole position from Gen4. This category is a no-brainer—PCIe Gen4 is fast. Really fast. There's not much else to discuss. The transfer format has a theoretical maximum rate of 8GB/s, and while current drives don't hit those speeds, it's still easy to see a real-time transfer rate of 5GB/s. PCIe Gen3 drives tend to cap out at around 3.5GB/s or below, while SATA drives are lucky to hit 500MB/s. Paying anything less than top dollar for Gen3 SSDs is almost sure to result in speeds more in the realm of 2.5GB/s, too.

There's more to performance than transfer speeds, though; Gen4 drives heat up faster than Gen3 models, which heat up faster than SATA ones. As those temperatures rise, performance throttles, which makes transferring large files problematic without good cooling solutions. X570 motherboards come with dedicated fans to mitigate additional heat from Gen4 M.2s, which does help. SATA drives handle heat well and perform admirably across large transfers, but when their base transfer speed is so low to begin with, it's just not enough to fully justify SATA drives over M.2. We'll give the win to Gen4 PCIe M.2 for those blinding fast speeds, but really it's whatever is best fit for purpose.

**Winner: PCIe Gen4 M.2**

## ROUND 3

### Compatibility

A while back, we would have given this to SATA any day of the week. It's practically impossible to buy a modern motherboard that doesn't support multiple SATA connections, and that has been the case for a while. However, M.2 2280 ports have now become more commonplace, making your choice of motherboard less of an issue if you're looking to install an M.2 drive. It's worth bearing in mind that if you're looking to use a PCIe Gen4 SSD, you're going to need an X570 motherboard and a third-generation Ryzen processor, which does limit your options quite severely.

Gen3 drives, on the other hand, have no such limitations with popular modern boards. As long as you're not going back more than a couple of generations, you can snag both Intel and AMD motherboards with at least one M.2 slot capable of supporting PCIe Gen3. SATA gets bonus points for usually supporting numerous channels, often up to six ports, even on cheaper boards, while M.2 ports take up more space on the board, so it's rare to see more than two or three. It's a small price to pay for not having to use cables, though, which gives the edge to the M.2 format—and in that arena, PCIe Gen3 is still king.

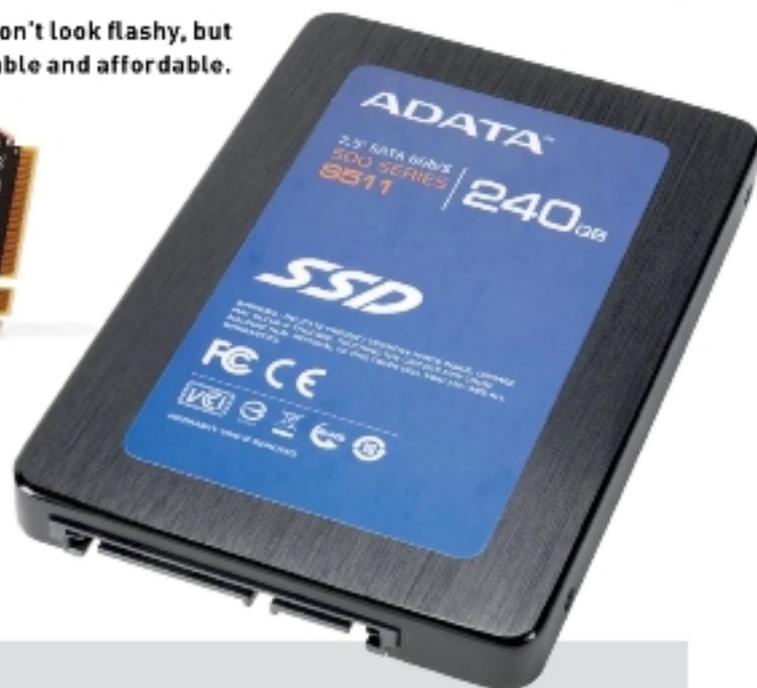
**Winner: PCIe Gen3 M.2**

# THE AD

There aren't many Gen4 M.2 drives available yet; this Gigabyte Aorus is one of the few.



SATA drives don't look flashy, but they're reliable and affordable.



Samsung is the king of Gen3 M.2 drives right now.

## ROUND 4

### Efficiency

Power-wise, SATA has a slight edge here. There's not much in it—broadly speaking, power consumption scales up as you increase both the capacity and transfer speed of the drive, but only by a small amount. A SATA and Gen3 M.2 drive of the same capacity might differ in power by a single watt or less, while a Gen4 M.2 drive would be perhaps 2W or 3W more. Gen4 drives have superior power management, though, consuming very little power when not actively reading or writing. As we mentioned earlier, SATA drives handle heat well (in part due to their larger design), but with the right precautions, M.2 temperatures should never reach a point where they're negatively impacting day-to-day use.

It's a close call here, especially considering that SSDs of all varieties never hit the power draw figures of components such as GPUs and coolers, which can be more than 10 times as demanding. All in all, it's a very tight decision and (were we to assign weightings to these categories) efficiency isn't the biggest factor you should consider while deciding between SATA and M.2, unless you're attempting to build an extremely low-power system. With that in mind, we'll give the win to SATA, but only by a hair.

**Winner: SATA SSD**

## ROUND 5

### Capacity

Although the technology behind M.2 SSD capacity is improving, SATA still has the edge. Architectural limitations mean that current commercial M.2 drives—both PCIe Gen3 and Gen4—are effectively limited to a maximum capacity of 2TB. Larger PCIe drives are available, but they either use an esoteric format other than M.2 2280, or simply cost an outlandish amount, such as Samsung's 4TB PM983, which will set you back \$3,500. Sure, 4TB SATA drives aren't cheap, but they don't hit quadruple digits, making them much more accessible for the average consumer. Even larger SATA SSDs are available, but again, these aren't just something you can order on Amazon. It's worth considering the lower bounds here, too.

SATA drives are typically available in a wider range of capacities, and are cheaper at the lower end, too, although very small capacity drives are swiftly becoming an antiquity that no manufacturer is willing to produce. Still, a good-quality 120GB Gen3 M.2 is still going to run your wallet more than a 120GB SATA SSD, and Gen4 M.2s are currently only available as small as 500GB, which obviously has a stern price tag attached. In the battle for capacity, then, the victor is SATA.

**Winner: SATA SSD**

## And the Winner Is...

There are arguments to be made for each of these types of SSD. PCIe Gen4 provides blisteringly fast transfer speeds; SATA is cheap and effective; PCIe Gen3 occupies a comfortable middle ground. For that reason, Gen3 is the winner; it's the best of both worlds, providing speedy transfers at a reasonable price. Unless you're really on a budget, it's hard not to recommend M.2 drives to anyone planning an upgrade or an entirely new build. Gen3 M.2 SSDs are plentiful, not too expensive, and provide a huge performance boost over SATA drives.

Really, though, it's important to remember that all of these drives are perfectly viable choices depending on the purpose of your build. A small but super-speedy M.2 SSD can be perfect as a boot and primary storage drive, with two cheaper and larger SATA drives supporting it for additional storage. SATA drives are perfect for budget builds; PCIe Gen4 is ideal for workstation systems, where the user will be moving numerous files around on the regular. We'll give the win to Gen3 M.2 here, but the real winner is the one that you decide in your heart. (And in your brain.) 🔄

# DOCTOR

THIS MONTH THE DOCTOR TACKLES...

- > Surface Solutions
- > Inactive Hyperlinks
- > Handling Handoffs

## Surface RT Updates

I have a Surface RT, the old-school first edition, and hope you can advise me. I had ditched this archaic device for an Android tablet, which worked great, but alas was taken by an ex-girlfriend in a situation I also wish you could help me with (but will spare you). I write a lot, so brought the RT back to life—or so I thought. The problems are beyond belief. I've checked with Microsoft, Googled until my brain hurts, and this Surface RT is just unable to update properly or use the Store, which makes it further obsolete because I can't receive anything considered new (8.1 RT), thanks to an error regarding network speed. I've checked my Internet connection, and it's fine. I've attempted to go Ethernet, but there's nothing compatible. I even bought Microsoft's USB Ethernet adapter. Am I out of luck, or is it possible to revive an RT?

—Ryan Anthony

**THE DOCTOR RESPONDS:** All these problems are related, and you can resolve them by resetting your Surface, but not in the usual fashion—it wouldn't be Microsoft if there weren't some tortuous

**Believe it or not, the original, neglected Surface RT is officially supported until 2023.**



process to jump through. Thankfully, this convoluted path has been trodden by others, so will hopefully work.

Connect your RT to the power, then restore it using its recovery image—head to <https://support.microsoft.com/surfacerecoveryimage> to obtain this if necessary, and follow the guide at <https://support.microsoft.com/kb/4037673/> if you need it. Once done, go through the setup process, then go to the Power and Sleep settings to set "When plugged in" to "Never"—given the amount of time you'll have to wait for some of these updates to appear, it's critical you prevent your Surface from sleeping.

Next, head to Windows Update to perform the first (of

many) searches for updates. This first one could take the best part of an hour, but once done, should bring up a firmware update. Install this, restart and return to Windows Update. Ignore the fact it's found more updates—a critical one (kb2919355) is missing, so search again, and it should appear in the list, at which point install the updates and restart when prompted.

The update process hangs if you try it again at this point. You need to download two updates manually: kb3173424 and kb3172614. Save these to your Surface's "Downloads" folder, then switch off the Wi-Fi (yes, really), manually install kb3173424, reboot,

manually install the second update, then reboot again.

Now turn your Wi-Fi back on and—if all has gone to plan—after around 10 minutes or so, a slew of new updates will come your way. Once in place, you'll hopefully discover your network issues are resolved, while you should also now be able to access the Windows Store to install apps.

## Can't Open Hyperlinks

I have a challenge for you: At some point in the recent past, possibly when Chrome was installed behind my back, I lost the ability to directly click on a hyperlink. I have tried everything I could find on the Internet to fix this headache, including reinstalling and removing IE, Chrome, Firefox (my browser of choice), and Edge, all to no avail. The error occurs if I click on a hyperlink, regardless of the program. The error message is: "This operation has been canceled due to restrictions in effect on this computer. Please contact your system administrator."

I may need to use Alex's "Windows Made Fresh" article, but I am very concerned about what will stop working.... —Will Collier

© MICROSOFT

submit your questions to: [doctor@maximumpc.com](mailto:doctor@maximumpc.com)

#### THE DOCTOR RESPONDS:

Without knowing exactly what you've done, the Doc can only offer generic advice. Three steps to follow are reassigning the default browser in Windows—switching from one browser to another and back again via “Start → Apps → Default apps.” If this fails, try clicking “Set defaults by app” from the same dialog—when the list appears, select “Firefox,” and click “Manage” to try to reassociate HTTP and HTTPS entries manually.

You may also need to reset your browser to its defaults—open Firefox, click the menu button, choose “Help → Troubleshooting Information,” and click “Refresh Firefox.”

Still no luck? Time to dive into the registry: press Win-R, type “regedit” and hit Enter, then navigate to HKEY\_CURRENT\_USER\Software\Classes\. Within this key locate the .htm key and change its (Default) value to htmlfile. Repeat for .html and also .shtml, .xht, .xhtml, and .xhtml keys, if they exist. Restart your PC, and you should find that you can select Firefox as the default browser again.

#### Restore Single Screen

I've been using my Windows 10 PC for years with two monitors: one in landscape mode, the secondary one in portrait. I'm going to set up my new PC in the same way, but I'll keep the old PC—mainly for backup—but use only one monitor with it.

I know how to set up my new PC for use with two displays, but have a question about switching from a two-monitor display to a single one: Will I need to first reconfigure programs that have been set up for two monitors so they'll work correctly with a single display, or will everything revert automatically to the one display? —**Doug Schafer**

**THE DOCTOR RESPONDS:** You shouldn't need to reconfigure anything, Doug. Programs should only use both displays

when they're both on and detected; if one is disabled, all the elements on the secondary display are dumped back on the main display at the edge of the screen that adjoins the display you disabled. You can then reconfigure them for the single display.

The best thing to do is configure this while both displays are still connected: Open “Settings → System → Display” on your old PC, scroll down to the “Multiple displays” drop-down, and set it to “Show only on x” (“x” being the number allocated to your display according to the “Rearrange your displays” image). Click “Apply” and you'll instantly see what happens when you revert to a single display—you can then confirm this before tidying things up, and disconnecting the second display.

#### Wi-Fi Mesh Issues

I have an Orbi RBR50 hub with satellite. The primary coverage is 2,000 square feet. The overall saturation and signal strength is good. Stationary devices, like the TVs, sync and maintain good strength and speed. My issue is with mobile devices—Galaxy phones and Kindle devices. As you move through the house, there is either no handoff from one Orbi device to the other, or the handoff is so slow that you drop the Wi-Fi connection—which means a poor or lost call when I'm using Verizon Wi-Fi Calling.

Netgear blames my phones, but they worked with a \$40 network extender. What am I missing? —**Greg Grubb**

#### THE DOCTOR RESPONDS:

The Orbi was one of the first mesh devices released, and it seems to have suffered more than its fair share of teething problems. These early implementations rely on client devices to detect which signal is strongest and switch accordingly—Galaxy devices are notorious for stubbornly holding on to weaker signals. Later mesh systems include a management controller whose job it is to scan for mobile devices to make sure they're using the strongest signal, forcing them to switch when they don't play ball.

The only real workaround is to reduce your Orbi router's transmit power, so there's a greater variance in signal strength between it and the satellite, making it easier for mobile devices to detect when to switch. You can both verify and alter the Orbi's signal strength via its web client by navigating to “Advanced Setup → Wireless Settings”—it can be reduced to as little as 25 percent. If this isn't doesn't work, place a physical barrier around the router (put it in a drawer, for example).

Still no luck? Given you've only got one satellite, it might be time to ditch it in favor of bringing that Wi-Fi extender back to life. It's cheaper than ditching the entire system for a newer mesh system.

#### Repetitive Repairs

I purchased a refurb Lenovo L530 with Windows 10 Pro and SSD about a year ago. It's recently started bringing up the automatic repair on every boot and restart, starting normally once it's run its diagnostics. I've repaired the MBR, run fixboot and sfc /scannow—any fix I can find on the Internet—without success. I tried rolling back with System Restore and using reset. Problem remains. I've even tried a different SSD (transferring my Windows install via True Image)—same result. An in-place reinstall keeping existing apps works for a couple of reboots, then it comes back. My last option is to wipe the drive and reinstall everything. Do you think this will work? Anything else I could try? —**Lionel Grindstaff**

#### THE DOCTOR RESPONDS:

Given you've gone to the trouble of taking a drive image, the Doc recommends going for the nuclear option, wiping the drive, and reinstalling Windows from scratch following our guide in the October 2019 issue. You've pretty much eliminated all other potential causes—make sure you create installation media from the Windows Media Creation tool, as this installs a clean version of the latest build of Windows 10. If the problem persists, you have a deeper—presumably hardware-related—issue, but you'll be able to roll back your system using the drive image.

If—as we suspect—the problem is fixed, but it later comes back, then you'll need to make a note of what software or hardware you recently reinstalled; focus your search for solutions around Googling “automatic repair” in conjunction with the program or device in question, in case it's a known problem, and you can pinpoint an exact cause to resolve the issue. Failing that, remove the culprit to hopefully eliminate the problem. ☺



# LIQUID-COOLED MONSTER

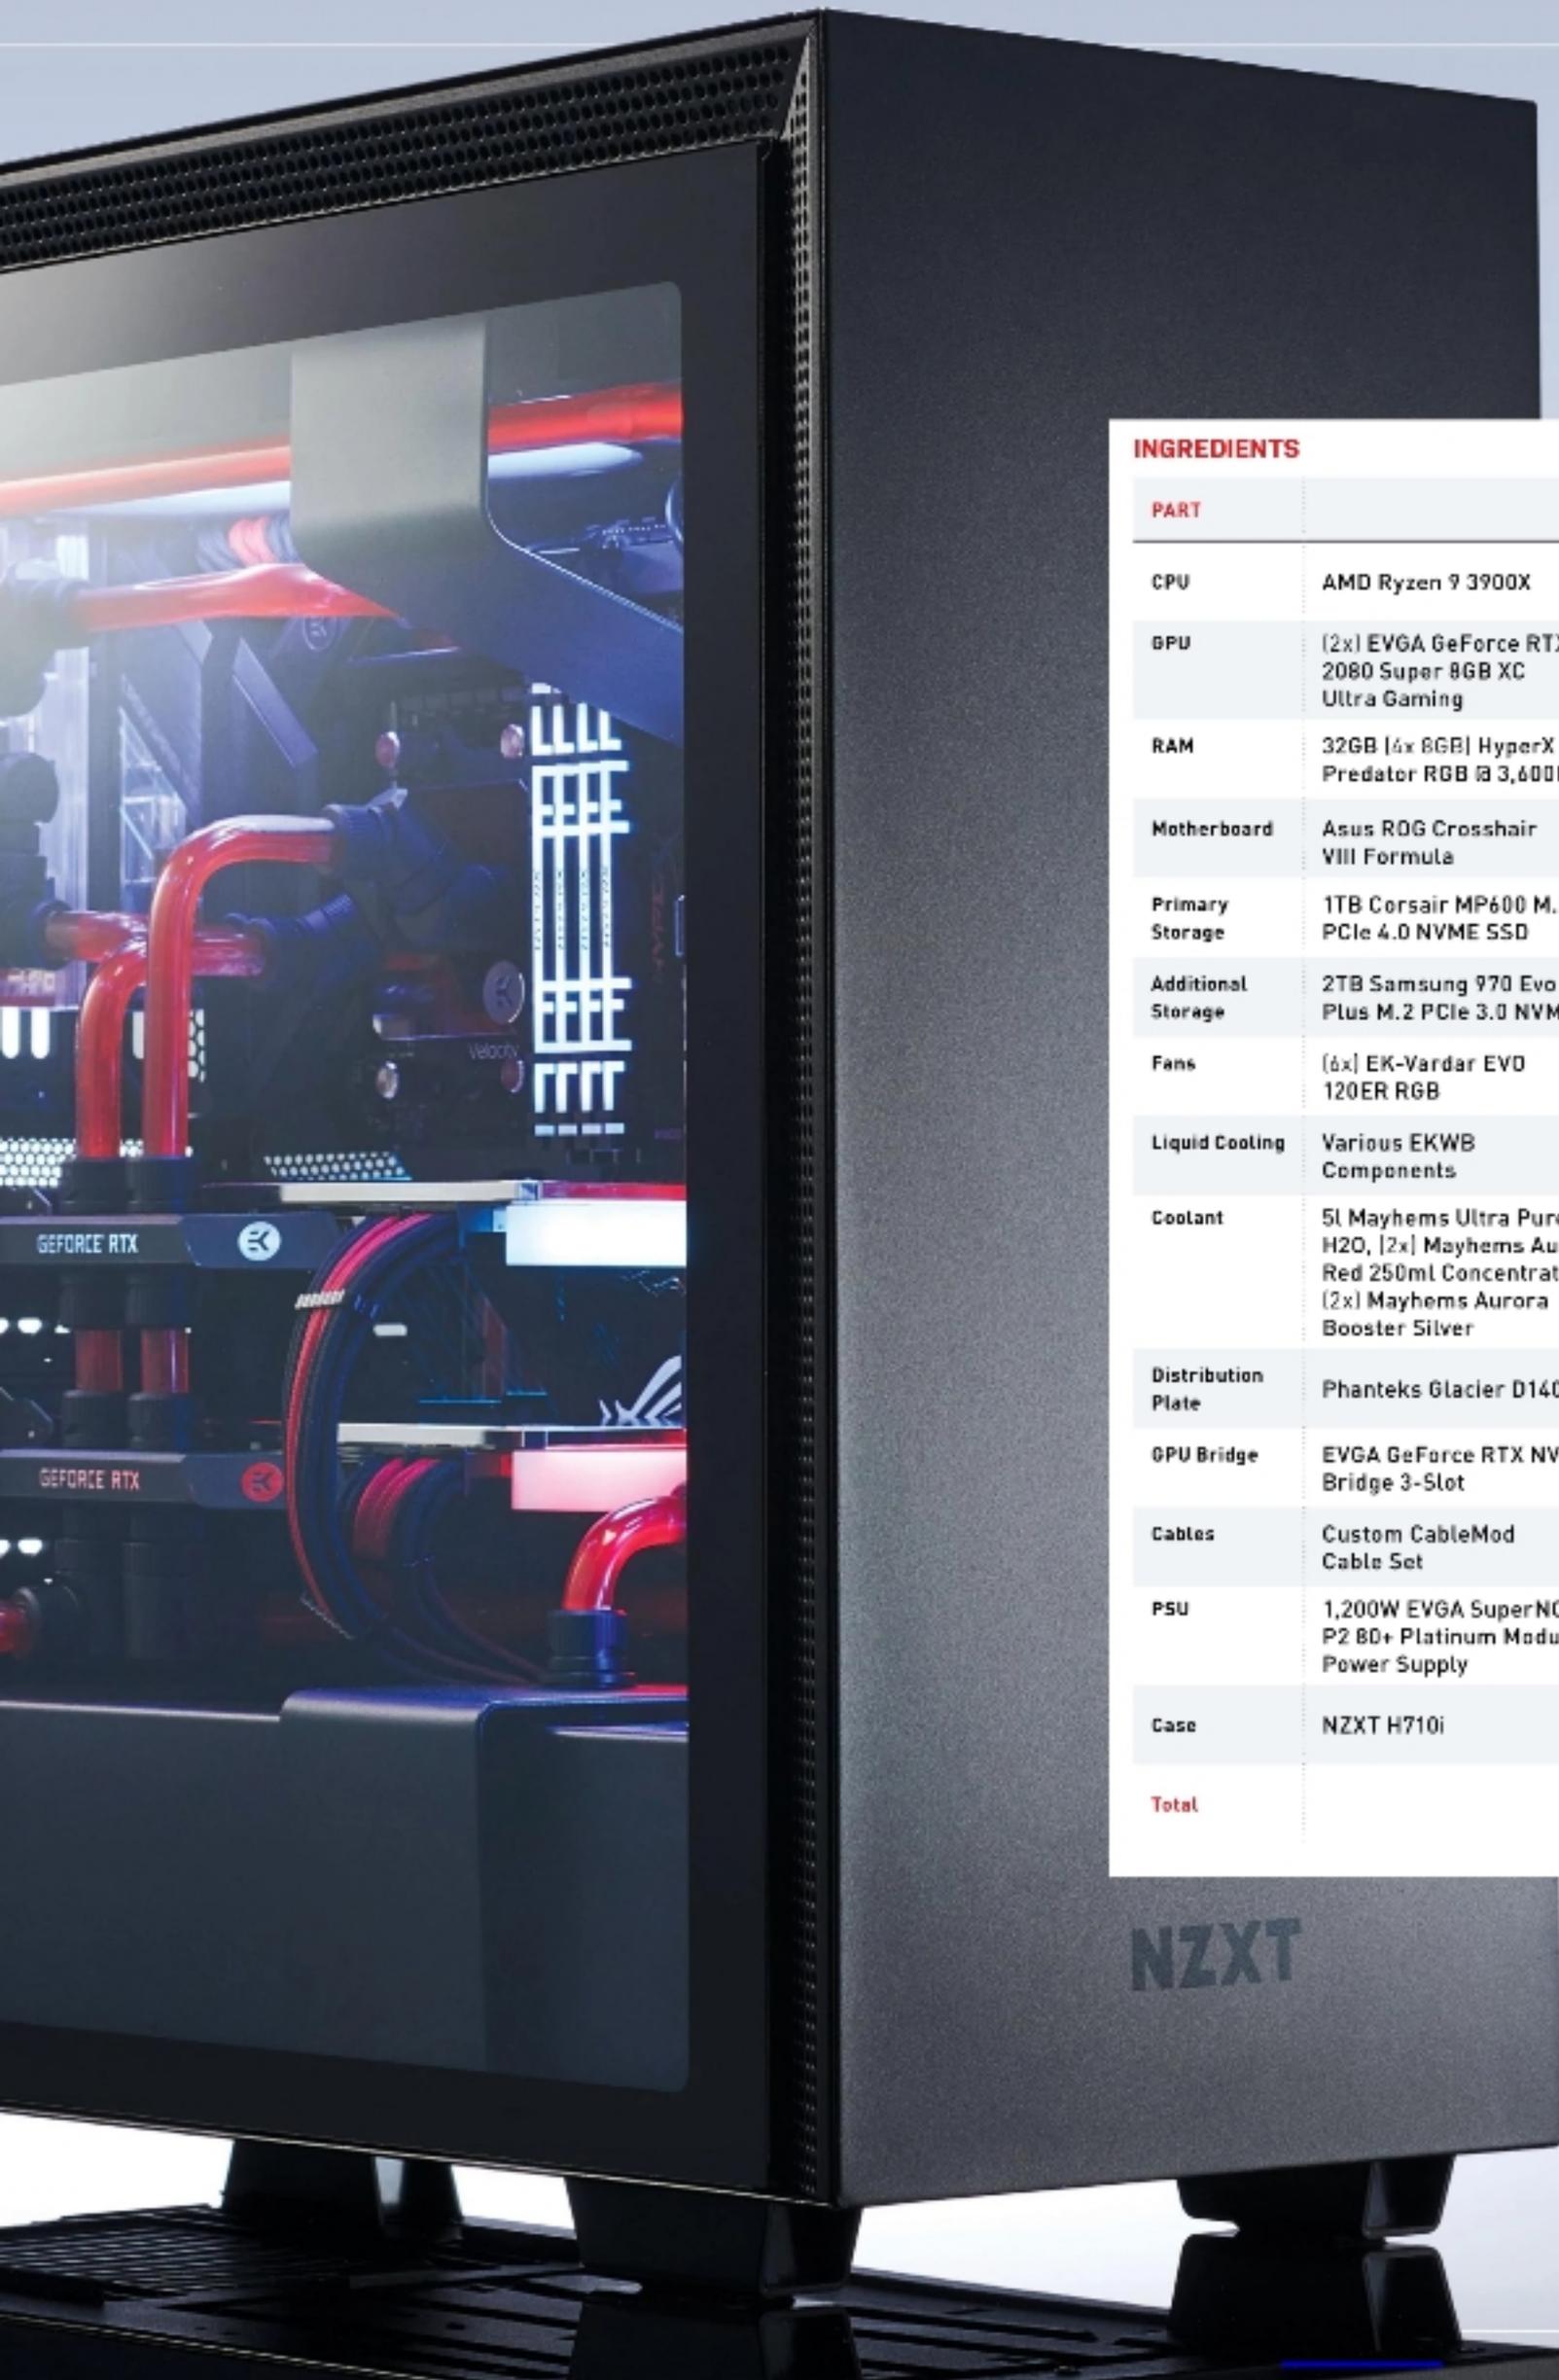
## THE \$6,000 RYZEN 4K SLI BUILD

**LIQUID-COOLING IS A FICKLE THING.** It's complex, intricate, and filled with far deeper considerations than any standard PC build. There's a finesse and beauty to it, and a risk involved. The very best systems out there incorporate it into their wonderful forms, yet the ideas and dreams of many a master builder often don't come to fruition, at least not how they were first imagined. If you're not committed to mastering this powerful art, and to overcoming the obstacles each build throws your way, it'll shake you off just as sure as the turning of the world.

This build has been a real challenge from start to finish. From the sponsor list, through the design phase and build process, to its real-world performance, every single element has come with its own set of defeats and victories. And yet, it is easily one of the best systems we've ever produced. Taking advantage of one of the latest Ryzen 3000 chips, and combining it with not one but two SLIed EVGA RTX 2080 Super GPUs, it was designed to truly push the limits of what you can achieve inside the impeccably orchestrated NZXT H710i case. It's a machine that asks questions: Are two 2080 Supers more cost-effective than one 2080 Ti? How much performance can you get from the very best Ryzen chip, in one of the very best Ryzen boards, when both are liquid cooled? Is the cost truly worth it in such an AIO-centric chassis? Hopefully, we can shed some light on all of those questions today.

Build aside, we are often asked whether liquid cooling is worth it. And the answer is that it depends on where you place your value. Financially, does it make sense? Absolutely not. It's the hypercar of the PC world—in reality, you could get just as much performance for a fraction of the cost by investing in some good fans and a chunky AIO. But artistically? It's beyond all measure, and the joy and pride you get from completing a build that is unique, that is truly yours, and that performs as magnificently as one like this does is priceless.

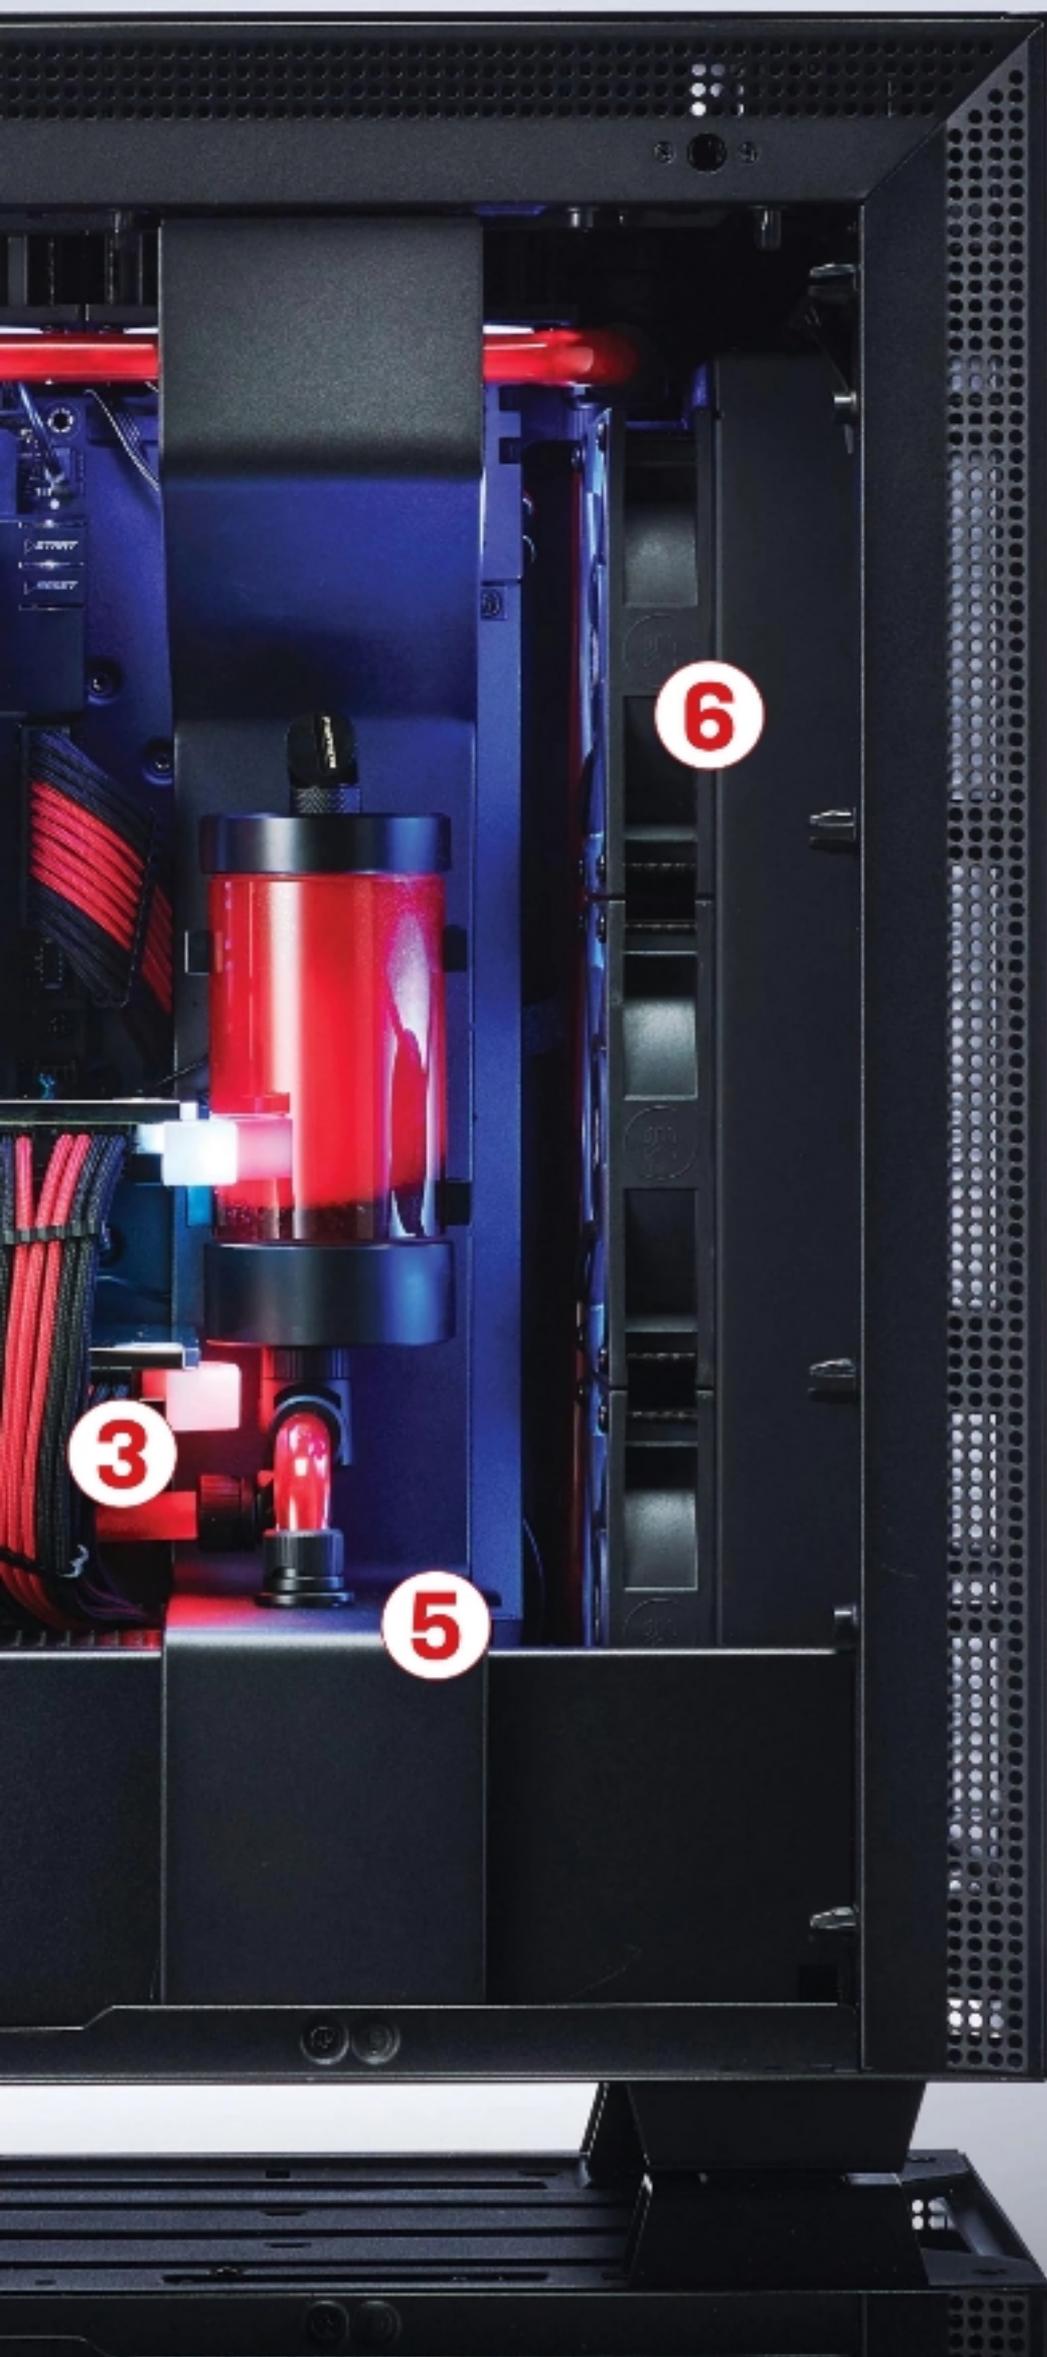




### INGREDIENTS

PART		STREET PRICE
CPU	AMD Ryzen 9 3900X	\$500
GPU	(2x) EVGA GeForce RTX 2080 Super 8GB XC Ultra Gaming	\$1,468
RAM	32GB (4x 8GB) HyperX Predator RGB @ 3,600MT/s	\$358
Motherboard	Asus ROG Crosshair VIII Formula	\$646
Primary Storage	1TB Corsair MP600 M.2 PCIe 4.0 NVME SSD	\$190
Additional Storage	2TB Samsung 970 Evo Plus M.2 PCIe 3.0 NVME	\$436
Fans	(6x) EK-Vardar EVO 120ER RGB	\$150
Liquid Cooling	Various EKWB Components	\$1,510
Coolant	5l Mayhems Ultra Pure H2O, (2x) Mayhems Aurora Red 250ml Concentrate, (2x) Mayhems Aurora Booster Silver	\$50
Distribution Plate	Phanteks Glacier D140	\$100
GPU Bridge	EVGA GeForce RTX NVLink Bridge 3-Slot	\$99
Cables	Custom CableMod Cable Set	\$271
PSU	1,200W EVGA SuperNOVA P2 80+ Platinum Modular Power Supply	\$238
Case	NZXT H710i	\$200
<b>Total</b>		<b>\$6,216</b>

PRICES CORRECT AT THE TIME OF PRINTING



### LIQUID COOLING INGREDIENTS

PART		STREET PRICE
CPU Block	EK-Velocity Strike RGB (Black Nickel & Black)	\$150
GPU Blocks	(2x) EK-Vector RTX 2080 RGB (Nickel & Acetal)	\$248
GPU Backplates	(2x) EK-Vector RTX Backplate (Nickel)	\$100
Reservoir 1	EK-RES X3 250 Lite	\$60
Reservoir 2	EK-RES X3 150 Lite	\$53
Pump	EK-XTOP Revo D5 PWM (Incl. sleeved pump)	\$140
Radiators	(2x) EK-CoolStream PE 360 (Triple)	\$176
Tubing	(6x) EK-HD Tube 10/12mm 500mm (2 pieces)	\$48
Standard Fittings	(28x) EK-Torque HTC-12 (Black)	\$196
45-Degree Fittings	(10x) EK-Torque Angled 45-Degree (Black)	\$100
90-Degree Fittings	(14x) EK-Torque Angled 90-Degree (Black)	\$140
Pass-Through Fittings	EK-AF Pass-Through G1/4 (Black)	\$8
Extender Fitting 1	(4x) EK-AF Extender 8mm M-F G1/4 (Black)	\$12
Extender Fitting 2	(2x) EK-AF Extender 20mm M-F G1/4 (Black)	\$10
Extender Fitting 3	(2x) EK-AF Extender 30mm M-F G1/4 (Black)	\$12
Extender Fitting 4	(2x) EK-Extender G1/4 Socket (Black)	\$10
Extender Fitting 5	EK-AF Extender Rotary M-M G1/4 (Black)	\$6
Extender Fitting 6	EK-AF Extender Rotary M-F G1/4 (Black)	\$6
Plugs	(4x) EK-CSQ Plug G1/4 with EK-Badge (Black)	\$20
Additional Fittings	Alphacool 90-Degree Rotary Fitting	\$15
<b>Total</b>		<b>\$1,510</b>

PRICES CORRECT AT THE TIME OF PRINTING

## HARDWARE COST: \$4,036



**CPU**  
AMD RYZEN 9  
3900X **\$500**

It might not be super-hot for high-end overclocking, but this 12-core processor will handle just about any game we throw at it, even CPU-bound large-scale strategy titles. With a base clock of 3.8GHz, there's not much this chip can't do, whether it's 8K editing or 4K gaming. The reported boost clock of 4.6GHz is attainable in individual cores, and it performs admirably when using auto-overclocking features. The support for PCIe Gen4 means that we've got room for a PCIe Gen4 M.2 drive running at maximum speed, too. All in all, this is one of the best desktop CPUs on the market right now.



**RAM**  
32GB (4x 8GB) HYPERX PREDATOR  
RGB @ 3,600MT/S **\$358**

HyperX makes some fantastic memory kits for high-end builds, and Ryzen thrives on high-speed RAM. The faster you can get that memory clock rate, the faster the Infinity Fabric

interconnect intertwining all those core complexes operates, meaning higher multicore performance.

With third-gen Ryzen, the sweet spot lies at 3,600MT/s, and we're taking advantage of that

with this 32GB kit of HyperX Predator RGB. Throw in some flashy lighting, a CAS latency of 17, and this finely-tuned kit of sticks will be more than enough for all our 4K gaming aspirations.

**CASE**  
NZXT H710i **\$200**

An awesome case for an awesome build. The H710i is the crown jewel of NZXT's mid-tower range, with USB-C support, cable management channels, and integrated RGB lighting. More specifically, there are two LED strips, which can be controlled via a smart hub hidden behind the cable bar. The tempered glass window uses a single thumbscrew for removal, making opening the case for adjustments easy. NZXT's iconic H-series cable management bar makes a return, too, which is handy. There's also a useful removable bracket for fitting fans or radiators, simplifying our cooling setup.





**MOTHERBOARD**  
**ASUS ROG CROSSHAIR VIII FORMULA \$646**

It's a bit on the pricey side, but everything about the Crosshair VIII Formula screams "quality." From the sleek RGB lighting to the Wi-Fi 6 support, and comprehensive onboard heat management solutions, this motherboard will pick up and carry your system through even the most nail-biting of battle royales. There's support for PCIe Gen4 SSDs, of course, but also gold-plated audio jacks, a thermal sensor connector, five-way optimization that improves overclocking performance, and too many other nifty features to list here. It's a powerhouse board, the perfect base for this build, keeping bottlenecks off the table.

**PSU**  
**1,200W EVGA SUPERNOVA P2 80+ PLATINUM MODULAR POWER SUPPLY \$238**

A modular PSU with a mildly hilarious number of ports—many of which we'll use for this build. With 92 percent Platinum-rated efficiency and a huge 1,200W capacity, the SuperNOVA P2 eliminates power bottlenecks just by existing. Two GPUs would put a strain on any system's power management, but not here; with this much power, we have ample room for overclocking every part of this system. It's a hefty unit, at over 9lb, and is optimized to support linked Nvidia GPUs, too. It also comes with an unbeatable 10-year warranty, so there's no need to worry about it blowing up and leaving you without power.



**SSD 1**  
**1TB CORSAIR MP600 M.2 PCIe 4.0 NVME SSD \$190**

What is an X570 mobo without a PCIe 4.0 M.2 drive? It's an old Scottish proverb, we swear. Really, though, this system would be incomplete without a speedy fourth-gen M.2 SSD as its primary drive. These PCIe 4.0 drives are scarce right now—at the time of writing, only 10 are available, spread across just four manufacturers—but performance is broadly similar, and this Corsair MP600 is the best value for a 1TB Gen4 M.2, at less than \$200. Expect awesome transfer speeds in the realm of 5GB/s.

**SSD 2**  
**2TB SAMSUNG 970 EVO PLUS M.2 PCIe 3.0 NVME SSD \$436**

1TB of storage wouldn't be enough for this build. The Crosshair VIII Formula has two M.2 slots; we're filling both. For our second drive, we've got a high-quality Samsung model, the 970 Evo Plus, bringing an extra 2TB of storage. Samsung's drives aren't cheap, but the performance here is guaranteed, with some of the fastest transfer speeds you're likely to see on any third-gen M.2 drive. Oddly enough, Samsung has yet to make the leap to PCIe 4.0, but that's OK—as this isn't our primary drive, we don't need the dizzying speeds of Gen4.



**GPU BRIDGE**  
EVGA GEFORCE RTX  
NVLINK BRIDGE 3-SLOT **\$99**

A simple piece of hardware that looks like something Batman would wear on his wrist while beating up nameless henchmen on the mean streets of Gotham. This is what gives our double-GPU setup its cadence; optimized for 4K gaming and ray tracing on Nvidia GPUs, EVGA's NVLink is one of the best GPU bridges around. Ours is the three-slot spaced model, although a four-slot version is also available for larger motherboards. There's some natty RGB lighting on it, too, if you're inclined to make your build more colorful.

**GPU**  
2x EVGA GEFORCE RTX 2080  
SUPER 8GB XC ULTRA GAMING **\$1,468**

That's right, we're using two GPUs again. Which is why we need the EVGA NVLink bridge (see left) to connect them. Looking at the cards themselves, we're pairing a couple of brand spanning new GeForce RTX 2080 Supers from EVGA; specifically, the XC Ultra Gaming model, with 8GB of GDDR6 VRAM, and a boost

clock of a stonking 1,845MHz. The original RTX 2080 was no slouch; the newer Super version ups the performance by a modest amount without damaging the price tag. There's full ray-tracing support, too, of course, with enough power to handle proper ray-traced reflections without tanking your frame rate.

**CABLES**  
CUSTOM CABLEMOD  
CABLE SET **\$271**

We're using custom cables for our power supply. Is that necessary? Maybe not. But these are braided, high-quality, and look better than the EVGA SuperNOVA's stock cables. Go to CableMod's website and you can build your own custom kit—in this case, a 24-pin ATX cable for mobo power, eight-pin and four-pin EPS cables for the CPU, and two bridged 8+6-pin connectors for our graphics cards. We've opted for flexible Paracord sleeves in red and black, with 600mm of length (700mm for the EPS cables).



**ACCESSORIES**  
COST: \$370



## COOLING COST: \$1,810

### FANS

6x EK-VARDAR  
EVO 120ER RGB **\$150**

Six 120mm fans for optimum case cooling, each providing high static pressure with double ball-bearing setups that provide a truly low-noise experience. With speed ranges of

500–2,200rpm and full PWM control, these are for high-performance systems. The RGB lighting is top-quality, too, with nine individual LEDs per fan—more lights than blades. Lastly, these fans

come with a unique smart-stop function, which fully switches off the fan when minimum fan curves are set below 25 percent and internal temperatures are sufficiently low, saving power and reducing noise.



### COOLANT

5L MAYHEMS  
ULTRA PURE H2O,  
2x MAYHEMS  
AURORA RED 250ML  
CONCENTRATE, 2x  
MAYHEMS AURORA  
BOOSTER SILVER **\$50**

We're using five liters of Mayhems' Ultra Pure H2O coolant to fill the tubes and reservoir, with its Aurora Red dye to color it. The Ultra Pure coolant lives up to its name, with a 10-step purification process and minimal electrical conductivity. Lastly, we're adding spice in the form of Aurora Booster Silver, which adds reflective particles to the coolant.

### LIQUID COOLING

VARIOUS EKWB COMPONENTS **\$1,510**

We'll get into the nitty-gritty of these parts during the build, but essentially we're using a full selection of EK components for our custom loop in this system. That means numerous 45- and 90-degree connectors, plugs, extenders, and a full meter of EK-HD tubing. Most importantly, there's the EK-XTOP Revo

D5—that's our full-sized PWM pump, charged with keeping the fluid in this cooling loop flowing. We've also got two EK-Vector waterblocks for our two GPUs, specifically the RTX 2080 RGB models; they're compatible with our 2080 Supers, and bring some additional lighting effects to this build.



### DISTRIBUTION PLATE

PHANTEKS  
GLACIER D140 **\$100**

Ordinary loops are for suckers. Designed to fit an ATX case with space for a 140mm rear fan, the Glacier D140 makes planning a custom cooling loop a cinch, thanks to dedicated ports for CPU and GPU cooling. Fitting around your rear I/O, it has plenty of air vents to ensure good case airflow, and uses tailored rubber sealing with excellent durability against temperature and stress. It also looks cool, with a mirrored backplate and integrated addressable RGB lighting.

## TOTAL OVERALL COST: \$6,216

# BUILDING THE MONSTER

A STEP-BY-STEP BUILD-IT GUIDE, DIVING INTO THE COMPLEXITIES OF THIS AQUATIC BEAST

LENGTH OF TIME: 20-30 HOURS

LEVEL OF DIFFICULTY: HARD

## 2 SEAT THE CPU BLOCK

**WE KNEW WE WANTED** to go with a mirror-finish EK block for this one, and fortunately the company produces a variant of its new Velocity line with a diagonal finish that matches our Asus motherboard styling. The only downside is that it comes configured for an Intel setup as standard. There is an additional Ryzen bracket included, but there are no instructions provided or available online on how to actually swap it in. To do that, you have to disassemble the block (there are four Allen screws on the bottom), swap out the original bracket, then secure it all down again, making sure you don't forget the rubber o-ring. A real nail-biting moment.



## 3 INSTALL STORAGE

**FOR THIS BUILD**, we're sticking with M.2 storage. Not only is it the fastest in the business, but it's also really convenient when you need to save space and reduce cable mess in a build. We've gone for a Corsair MP600 Force 1TB PCIe 4.0 drive and a Samsung 970 Evo Plus 2TB PCIe 3.0 drive, both hidden away under a heatsink on the Formula. You have to strip the Corsair drive out of its massive incorporated heatsink for this by popping two clips on the side, but once that's done, you're good to install it into your motherboard. Don't worry, you can always retroactively reinstall it into its heatsink after.



## 1 ADD GPU BLOCKS

**IT'S ALWAYS SCARY** taking a graphics card apart, because you're messing around with something that's insanely expensive. On the flip side, liquid-cooling anything below an RTX 2080 or equivalent tends to be difficult to justify from a cost-to-performance perspective, so it's pretty much always the more expensive cards that you strip apart. In our case, it was two RTX 2080 Supers, which were provided by EVGA. We've combined them with two of EK's Vector RTX waterblocks, along with some mirror-finished backplates to help both dissipate heat and throw light around the system.



## 4 FIRST INSTALL

**SO, THIS IS** our first look at how everything's going to fit inside the case, and it's tight. We originally intended to run a 250ml reservoir horizontally below the second GPU—however, due to incompatible tolerances (read: It didn't fit), we had to opt for an alternate solution. Fortunately, we'd requested two different sizes of reservoir for just this situation.



## 6 FRONT RADIATOR COMPATIBILITY

**ON TO THE FRONT RADIATOR**, and the compatibility issue here comes from your tubing runs, and how you install your fans. If you're utilizing two 360mm radiators, as we are with our loop, there's only one way you can run your fans—with them on the inside. This gives you a little breathing room when it comes to connecting the radiator to other components. We also recommend keeping the G1/4-inch ports for your front radiator at the bottom of the case. Note that we have removed the cable bar entirely here, so we can get better access to the front of the case.



## 5 RADIATOR COMPATIBILITY

**IF YOU PLAN** to build a similar system in this chassis, beware of your top rad thickness. We're using one of EK's PE 360mm radiators; it's about 1.5 inches thick, but the top panel comes into contact with it at the rear of the case when installed. There is a way around this: You'll note the top panel is constructed of a metal sheet and a plastic sheet stuck together. Carefully unscrew the plastic sheet, pull it apart, cut the plastic layer in half, reattach it to the metal, and the top panel will sit flush. Problem with that is you'll achieve zero airflow (although it looks great for a photoshoot). What you really need to do is install a skinny radiator in the roof. We recommend one of EK's SE 360mm radiators, as they come in at 1.1 inches, giving you some room for exhaust. Or, if you're handy with a CNC machine, you could cut some ventilation holes in the top panel.



## 7 A TIGHT SQUEEZE

**FOR OUR 150ML RESERVOIR**, we're mounting it directly into the cable bar. However, it's a tight squeeze around our two GPU blocks. The cutout in the cable bar is specifically designed for reservoir mounting, and it works well. It's also worth noting that we eventually ended up changing the top of the reservoir, because the full-height version, complete with downpipe, that you see here is too big for an angled 90-degree fill port to be added without conflicting with the cable bar, and as we don't need the down tube, due to our loop configuration, it's no biggy swapping it for a lightweight one.



## 4 FIRST INSTALL

**SO, THIS IS** our first look at how everything's going to fit inside the case, and it's tight. We originally intended to run a 250ml reservoir horizontally below the second GPU—however, due to incompatible tolerances (read: It didn't fit), we had to opt for an alternate solution. Fortunately, we'd requested two different sizes of reservoir for just this situation.



## 6 FRONT RADIATOR COMPATIBILITY

**ON TO THE FRONT RADIATOR**, and the compatibility issue here comes from your tubing runs, and how you install your fans. If you're utilizing two 360mm radiators, as we are with our loop, there's only one way you can run your fans—with them on the inside. This gives you a little breathing room when it comes to connecting the radiator to other components. We also recommend keeping the G1/4-inch ports for your front radiator at the bottom of the case. Note that we have removed the cable bar entirely here, so we can get better access to the front of the case.



## 5 RADIATOR COMPATIBILITY

**IF YOU PLAN** to build a similar system in this chassis, beware of your top rad thickness. We're using one of EK's PE 360mm radiators; it's about 1.5 inches thick, but the top panel comes into contact with it at the rear of the case when installed. There is a way around this: You'll note the top panel is constructed of a metal sheet and a plastic sheet stuck together. Carefully unscrew the plastic sheet, pull it apart, cut the plastic layer in half, reattach it to the metal, and the top panel will sit flush. Problem with that is you'll achieve zero airflow (although it looks great for a photoshoot). What you really need to do is install a skinny radiator in the roof. We recommend one of EK's SE 360mm radiators, as they come in at 1.1 inches, giving you some room for exhaust. Or, if you're handy with a CNC machine, you could cut some ventilation holes in the top panel.



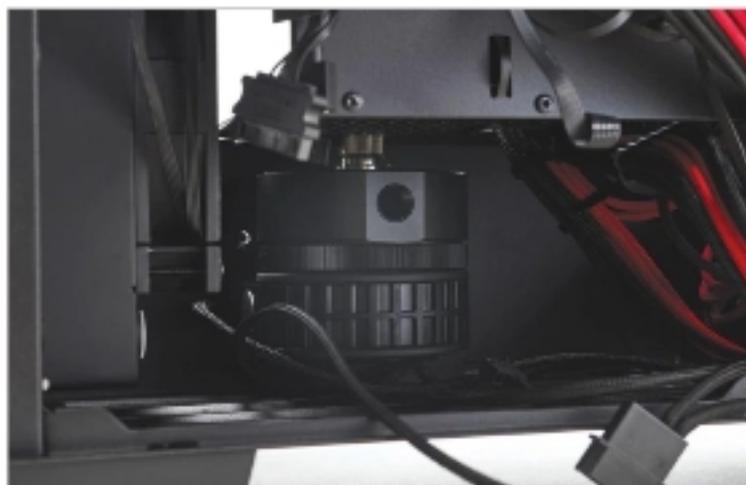
## 7 A TIGHT SQUEEZE

**FOR OUR 150ML RESERVOIR**, we're mounting it directly into the cable bar. However, it's a tight squeeze around our two GPU blocks. The cutout in the cable bar is specifically designed for reservoir mounting, and it works well. It's also worth noting that we eventually ended up changing the top of the reservoir, because the full-height version, complete with downpipe, that you see here is too big for an angled 90-degree fill port to be added without conflicting with the cable bar, and as we don't need the down tube, due to our loop configuration, it's no biggy swapping it for a lightweight one.



## 8 INSTALL THE PUMP

**WHEN IT CAME TO INSTALLING** our pump, we knew we needed it to connect to a pass-through that pierced both the cable bar and the PSU cover. We've used a male-to-male extension fitting to attach the inlet of the pump to the G1/4 pass-through, which will then connect to the reservoir above. We've also drilled four holes in the PSU cover, hidden behind the cable bar, and have mounted the pump to the cover. EK includes a handy markup template for you to use. It was tricky to install, but once the pump is in, it keeps the whole pump-res affair looking far tidier than running a combi unit. You can also see that we've pre-emptively installed the PSU and its cables as well, just to get them out of the way before we commence our tube-bending stage.



## 10 THINK ABOUT YOUR BENDS

**ALWAYS START OUT WITH** simple tubing runs first. You can't see it in this shot, but we've already done the straight tube running from the front radiator to the top one, and the ones surrounding the top of the distribution plate. It's always smart to focus on doing simple, clean 90-degree bends. There's a temptation when it comes to tube bending to overcomplicate things, and introduce multiple bends to one length of tubing, but we guarantee you that it'll look 100 times better if you keep it simple and take advantage of angled fittings instead. Here, you can see we've got our GPU bridge tubes in, our bottom distribution plate tube in, and we are test fitting a—clearly wrong—CPU tube as well.



## 9 VISUALIZE YOUR RUNS

**AT THIS POINT,** you're looking at the end of day one of our build photoshoot. We've now got the majority of our fittings installed and are looking for ways of connecting everything together. This might sound silly, but the best advice we can give you is to spend a lot of time in this phase simply staring at your case. Figure out where you want your tubes to go, what the best angles are, and where you'll need extensions. Then go and sleep on it. The setup we finally ended up with was very, very different from the one we had envisioned right at the beginning. In our case, we ended up dropping the distribution plate by a few millimeters, and changing out some of the fittings entirely to get better tubing runs.



## 11 CPU RUNS

**OUR FINAL SETUP** for the CPU is this parallel loop configuration. We've had to stagger the tubing runs in a step like this due to the new EK Torque fittings being, well, massive. We've effectively used four 45-degree fittings here to achieve this snakehead effect, to ensure they align perfectly with the ports on the topmost GPU, but they'll only twist so far each way until they hit each other and conflict, or stop the straight lines we're after. With one staggered below and one up top, it means we can still run two straight tubes, with single 90-degree bends without worry. In fact, we kinda like it.



## 12 PUMP-RAD MASTERY

**DESPITE BEING** perhaps the simplest tubing runs, this is still some of our proudest work. With a combination of extension and rotary fittings, we've managed to connect the furthestmost G1/4-inch port from the radiator to the pump using a single piece of straight acrylic. Then we've run another single straight piece of tubing direct from that bottom radiator (again with two 90s and an extension) straight up to the top radiator without a problem. Simple, but efficient.



## 14 NVLINK BRIDGE CONFLICTS

**WHEN RUNNING TWO CARDS** with waterblocks, finding an NVLink bridge that will fit is always going to be a problem. In our case, the ones we had conflicted with the ports on the blocks, so we've had to strip one down. Fortunately, with the EVGA bridge, it's a fairly painless procedure. Four small Phillips screws in the back enable you to remove the top plastic cover, showcasing a separate piece of PCB with the RGB LEDs on top, and a small ribbon cable on that connects the LEDs to the rest of the NVLink bridge. Disconnect the cable, apply a little heat via a heat gun to remove the majority of the glue, and voila—you've got a pretty snazzy-looking NVLink PCB bridge, working, and ready to go.



## 13 RESERVOIR PROWESS

**THE TWO RUNS** connecting the reservoir to its components are by far the most difficult. This 90-degree bend is very short; the closer you are to the bend radius with your cut, the harder it is to get the compression fitting cap on it, ready to secure to the fitting. We've had to use a 90-degree Alphacool rotary fitting in the bottom of the res at the back, because the EK Torque fittings are so large you can't fit two of them in the ports. Alphacool's are skinnier, so installing one hidden fitting in the back ensures we can get the tubing run we want, without compromising style or loop orientation. The tube connecting the bottom GPU to the reservoir in that Alphacool fitting is also tricky. It's a long 90-degree bend, which is fine, but as we have to secure it with the GPU in place, it's difficult to see, so tricky to know if it's secured properly. Double-check you've tightened your tubes in place as best you can before filling.



## 15 FILL IT UP!

**RIGHT THEN,** on to the filling. Remember to put plugs in all the G1/4-inch ports you're not using (bottom of the reservoir, on the pump, the distribution plate, GPUs, and so on), and you're good to go. We're using Mayhems Aurora Concentrate mixed with some Ultra Pure H2O and Aurora Silver Booster. When running a parallel loop like this, the procedure is much the same as usual. Fill the res, cycle the pump, and continue until full. The main difference is you're more likely to get air trapped in places you don't want it. This will move to the res over time, thanks to flow rate and pressure, but try to make sure you don't get air trapped in the CPU block if you can, because it dramatically affects temps and performance. In fact, once filled, physically tilting the case or shifting it on to its side and running the pump again (only when the system is full and sealed) is a good way of shifting the air to the reservoir.



# LIQUID-COOLED CONCLUSIONS

## Questions asked; questions answered

**UNSURPRISINGLY**, when it comes to liquid-cooling PCs, the more experience you have, the easier they become to build. However, even after spending three days together, and the system was finally primed, our photographer and system-builder turned to each other with a look of confusion: It's never this easy. Usually, if it can go wrong, it does. Yet we hit our deadline, nothing leaked, no tubes had to be replaced or reseated, and it all worked.

The system looks incredible, but as with all builds, it's not without its flaws. So, what would we change, what would we do differently if we had more time, and what will we do going forward?

### TUBING TWEAKS

The biggest issue from an aesthetic perspective is some of the tubing runs, most notably the one that connects the bottom of the distribution plate to the second GPU. It's a tight 90-degree bend, for sure, but we're not satisfied with the downward angle.

There are two ways we could fix this. We could add an 8mm extension spacer between the two 90-degree fittings to lower the height of the bottom-connecting 90-degree fitting, and therefore reduce the angle of tube, or we could lower the distribution plate by a fraction of an inch (which is relatively easy, as it's only mounted via four Allen screws in the back).

That last solution may seem extreme, especially as it could adjust the angle of the top three tubing runs as well. But we're confident we could get around that, as the top VRM run is comprised of two 45-degree fittings, which we can rotate to compensate, and the bottom VRM run also has tolerance to do the same.

The reason this is preferable over just installing an 8mm extension is because, as it's at an angle jutting out toward the

window, adding a spacer could cause the bottom fitting to come into contact with it, preventing us from being able to secure the window in place.

The bottom-most VRM run is also frustrating, as the tube is slightly too long where it connects to the 90-degree fitting, putting extra pressure on the fitting, and pushing it backward. It's secure for now, but over time it could form a leak if we're not careful. The solution would be to shorten that run by using a file to sand down the edges a touch. However, that might make it harder to secure the compression cap, as it has to go around a 90-degree bend first. It's not impossible, just frustrating, and harder to be certain whether the tube is seated correctly.

Lastly, the top tubing run connecting the uppermost port on the distribution plate to the uppermost rad could also do with a fresh piece of tubing and a new 90-degree bend, as it's ever so slightly short on the distribution plate side (we've compensated with an 8mm extension fitting here).

### HARDWARE CHANGES

Outside of those slight adjustments, there's not much we'd change from a hardware perspective. Our biggest issue is that top radiator: Even with our top panel's plastic interior cut in half, when the panel is secured on the chassis, it comes into direct contact with the edges of the radiator, so the only ventilation for all that exhaust air is a small slit at the back of the panel. That's not ideal.

The simplest solution would be to replace it with a skinnier SE edition 360mm radiator, similar in thickness to that found in many AIOs. That said, we quite like the extra surface area gained by the PE edition radiators, especially considering our hardware. For our

testing, we removed the top panel to provide sufficient exhaust to the fans beneath, but the long-term plan is replace the panel with a milled sheet of aluminum, with holes cut out for extra ventilation. Failing that, acrylic or wood may suffice.

### PRECISION PERFORMANCE

Now the design choices are out the way, it's time to talk numbers, and answer some of those questions we posed at the beginning of this article. The big talking point for us comes down to those twin SLIed (or, we guess, NVLinked) GPUs: the EVGA GeForce RTX 2080 Super XC Ultra Gaming cards.

To be honest, SLI has long been something we haven't really ever recommended, and when support for it started to wither away with the 10 series, we assumed it was all but dead. However, it seems that this is currently Nvidia's best kept secret. In 4K in particular, the uptick on running two RTX 2080 Supers, as opposed to a single RTX 2080 Ti, is astronomical.

### 4K & 1440P GAMING

In our 4K game benchmarks, we saw an average of 94fps in *Middle Earth: Shadow of War* and 82fps in *Total War: Warhammer II*, compared to 56fps and 44fps from our single liquid-cooled 2080 Ti build zero-point. That's an increase in performance of almost 68 percent in *Shadow of War* and 86 percent in *Total War*—again, that's compared to a single 2080 Ti, not a single 2080 Super.

Now, it would be unfair of us not to mention some cost figures here: Our SLI setup comes in at around \$1,410, or \$1,510 with the NVLink bridge; the 2080 Ti on its own comes in at \$1,219. So, for \$290 more, you're getting up to 86 percent more performance in 4K in certain titles. Of course, you have additional power draw to take into consideration (our system's total draw is around 724W under full load, versus 474W for the same rig and a single Ti), and you have the drawback (as shown by *Assassin's Creed*) of some games not taking advantage of the other GPU.

That last part is more interesting than you might think. If you run HWMonitor throughout your benchmarking, you'll spot how GPU loads vary across titles. In games that are properly optimized (*Total War: Warhammer II*, *Middle Earth: Shadow of War*, *Star Citizen* as examples), you'll usually see 95 percent load on both cards; on games that aren't (*World of Warcraft*, *Assassin's Creed*), the figures will fluctuate around 85–90 percent load on GPU 1, and 15–20 percent load on GPU 2. It's not all bad news in that scenario, as



## CPU BENCHMARKS

	ZERO-POINT	DREAM KILLER
Cinebench R15 Single (Index)	231	186 [-19%]
Cinebench R15 Multi (Index)	1,701	3,218 [89%]
TechARP x264 (fps)	37.61	64.81 [72%]

## STORAGE BENCHMARKS (MB/s)

	ZERO-POINT	DREAM KILLER
AIDA64 Memory Bandwidth	45,683	55,073 [21%]
AIDA64 Memory Latency (ns)	50.3	74.2 [48%]
SSD1: CrystalDiskMark Sequential Read	3,522	4,889 [39%]
SSD1: CrystalDiskMark Sequential Write	2,237	4,168 [86%]
SSD1: CrystalDiskMark 32-Bit 4K Read	406	568 [40%]
SSD1: CrystalDiskMark 32-Bit 4K Write	314	490 [56%]
SSD2: CrystalDiskMark Sequential Read	3,518	3,472 [-1%]
SSD2: CrystalDiskMark Sequential Write	3,244	3,133 [-3%]
SSD2: CrystalDiskMark 32-Bit 4K Read	425	427 [0%]
SSD2: CrystalDiskMark 32-Bit 4K Write	360	355 [-1%]



## 1440P GAMING BENCHMARKS (Avg fps)

	ZERO-POINT	DREAM KILLER
Total War: Warhammer II	75	116 [55%]
Middle Earth: Shadow of War	87	130 [49%]
Assassin's Creed Odyssey	54	60 [11%]

## 4K GAMING BENCHMARKS (Avg fps)

	ZERO-POINT	DREAM KILLER
Total War: Warhammer II	44	82 [86%]
Middle Earth: Shadow of War	56	94 [68%]
Assassin's Creed Odyssey	35	42 [20%]

## SYNTHETIC GAMING BENCHMARKS (Index)

	ZERO-POINT	DREAM KILLER
3DMark Fire Strike 1080p	22,503	33,086 [47%]
3DMark Fire Strike Ultra 4K	6,801	12,837 [89%]
3DMark Port Royal Ray Tracing	5,967	13,406 [125%]
3DMark Time Spy DX12	9,775	20,506 [110%]

## TEMPERATURE BENCHMARKS (°C)

	ZERO-POINT	DREAM KILLER
CPU: Prime 95 Blends Test	66	67 [2%]
CPU: Cinebench R15	75	74 [-1%]
GPU: 3DMark Fire Strike	58	56 [-3%]
GPU: 3DMark Port Royal	58	66 [14%]
CPU: 3DMark Fire Strike Combined	72	61 [-15%]
GPU: 3DMark Fire Strike Combined	59	59 [0%]

Our zero-point consists of a liquid-cooled and delidded Intel Core i7-8086K @ 5.3GHz, 32GB of DDR4 @ 3,200MT/s, an Asus ROG Maximus XI Formula, a liquid-cooled Asus GeForce RTX 2080 Ti, a Samsung 970 Pro 512GB M.2 PCIe SSD, and a Samsung 970 Evo Plus 1TB M.2 PCIe SSD. Highest recorded temperature shown; all games tested at highest possible profile.

the second card is not drawing anywhere near as much power, or running as hot, but it's still something to note. We did some testing in *WoW* to examine how GPU load and SLI were managed with different APIs. In DX12 mode, GPU 1 would take 85–90 percent load, and GPU 2 would take 15–20 percent. In DX11, bizarrely, the cards would split the workload, with both maxing out at around 45 percent load.

Back to frame rates for a second, and we will briefly say that at 1440p, you will notice less of an uptick in performance than at 4K, as frame rate starts to lean more on the processor. However, in games that support it, you're still looking at 50–60 percent performance bonus versus a single 2080 Ti.

### **RYZEN, PBO, & AUTO-OVERCLOCKING**

If you read last issue's review of the Asus Crosshair VIII Formula, you'll know we rated it very highly because of its fantastic PBO (Precision Boost Overdrive) performance. In short, there's software baked into the Ryzen 3000 series that measures various voltage and temperature parameters across the processor and VRMs. The cooler you can keep them, and the higher the quality of the power phase solution, the longer the processor remains at its highest

boost clock. Combine that with Ryzen's auto-overclocking feature, and with sufficient cooling, you'll get some serious performance out of it. As we did.

To give you an idea, most VRM solutions are designed to withstand temperatures up to 110 C; after that, they throttle the amount of voltage going to the processor. In our case, as our VRMs are liquid cooled, the temperature never rose above 37 C during any of our tests. Combine that with the relatively low CPU temps we achieved and the impressive auto-overclock on the turbo clocks, and performance was astronomical—in the long-term tests, in particular—compared to our stock review last month (which also topped charts back then). With the VRMs cooled at stock, the Crosshair and Ryzen 9 3900X scored 63.24fps in *x264*, 3,174 in Cinebench R15, and 75 seconds in Fry Render. Our liquid-cooled Formula achieved 64.81fps, 3,218 in Cinebench R15, and just 69 seconds in Fry Render, all at stock, without us tweaking the overclock anywhere.

Of course, it's a lot of cash to invest in a custom loop purely for some slightly faster processor performance, but if it comes as a side effect of cooling the CPU and GPU? Well, you might as well enjoy it, right? Of course, our 2080 Ti rig was powered by an Intel Core i7-8086K, and as such, with half

the cores and half the threads, that poor processor was murdered everywhere except the single-core figures, even when it was clocked to 5.3GHz—but that's not really a fair comparison, is it?

### **TEMPERATURES & TOLERANCES**

As for temperatures, this is where things get interesting. As our system is running a parallel loop, theoretically temperatures should be higher than with a serial loop, especially if you don't have sufficient flow. This system has its own personality, and it took us a while to fully understand it properly and get it to a state where it could benchmark most efficiently.

What we mean by that is that, thanks to that parallel loop (coolant can theoretically pass straight through the bottom GPU and out into the reservoir, without circulating around the top GPU or the CPU), there's a few requirements that need to be met before temperatures are what we would consider stable.

Firstly, running a parallel loop relies entirely on pressure within the system itself. If you've got a high enough flow rate from your pump, you'll find coolant will move throughout the system with no issue. As long as coolant is moving, heat is exhausted from the radiators. When we first started priming and setting the

# techradar.



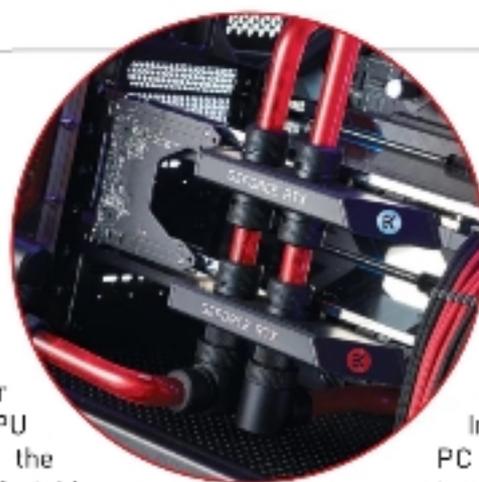
system up, we forgot to switch the fan header that the pump was running on to full speed, which meant that coolant was moving very slowly past the CPU and down back into the GPU. At full speed, however, thanks to the Mayhems Aurora coolant, you can clearly see that all the coolant is flowing quickly around these parts.

Secondly, air bubbles are a big deal. Unlike a serial loop, where bubbles naturally move to the reservoir over time, in a parallel loop, they have a habit of getting caught in the top of the tube connecting the CPU to the GPU—or, even worse, getting stuck in the CPU block, which you can't see into, of course. It took us a lot of time tipping the rig back and forth, and running it in various orientations, before we could move all of the air to the reservoir, and then replace it with coolant. Even after a week of continuous use, there are still small bubbles appearing in the topmost tube on the CPU. So, yeah, it definitely looks good, but if you're after pure performance and less fuss, it's a hell of a lot easier to run a serial loop.

Those issues aside, our temperature tests came out quite well. During our Prime95 Blends test, we maxed out at 67 C, Cinebench R15 saw that figure rise to 74 C, and in our graphics tests, the

cards performed admirably, with the highest card reporting 56 C in Fire Strike and 66 C in the ray-tracing powered Port Royal benchmark. In our combined CPU and GPU game test, we saw the CPU sitting at a comfortable 61 C, and the hottest GPU at 59 C. The second GPU (located at the bottom of the stack) usually ran between 9 and 13 degrees cooler than the primary card, although this is something we're used to seeing in SLI setups, regardless of cooling, so it's nothing we're particularly concerned about.

Compared to the Intel, RTX 2080 Ti rig, our build is considerably warmer across the board. However, we're cooling twice as many GPUs and twice as many cores, with the same amount of radiator surface area, so we'd expect temperatures to be higher. However, the RTX 2080 Ti is only 6–8 degrees cooler than our topmost card in this rig; the biggest difference is the Intel Core i7-8086K, registering 21 degrees cooler than its Ryzen 9 3900X competition. It's also worth noting that the Core i7-8086K was delidded, and the TIM



connecting the die to the heat spreader had been replaced with liquid metal, which typically reduces temps by 10–15 C.

#### **THE END?**

In our eyes, at least, this PC is quite possibly the best system we've ever built. The

arsenal of hardware within is what we consider to be some of the best high-end components you can get, weirdly mixing good value with an insane price tag. The cooling is unique, and despite our original concerns regarding a parallel setup, it performs admirably. SLIed 2080 Supers are nothing short of incredible in some games, and the processor also performs really well, even at stock.

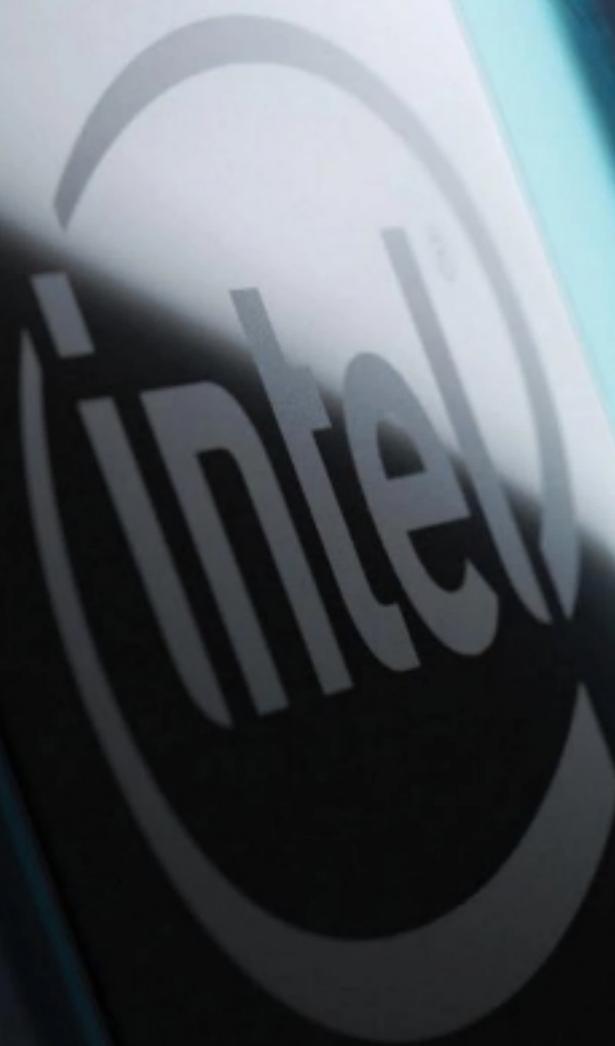
There are, of course, things we will probably change over the next few months, but as far as the build itself went, we faced few hurdles that really set us back, and despite the cramped conditions inside the H710i, it's a case that has some serious cooling potential.

Is it perfect? Absolutely not—it's insanely expensive for a start. But, damn it, this is *Maximum PC*, and building insane rigs is in our blood. ☺



The home of technology

**techradar.com**



You could argue that the 10th-gen Core is the successor to the Core 2 Duo from 2006.

# 10 IS THE MAGIC NUMBER

Intel's 10th-gen Core CPUs are here, and with them its 10nm production process. But, as **Jeremy Laird** found, the power of 10 is proving a little problematical

IT'S BEEN FIVE LONG years since Intel's very first 14nm CPUs came to market. That's worth repeating. It's taken fully five years for Intel to roll out the 10nm successor to its 14nm production node. This from the company that has led the industry in manufacturing integrated circuits for decades. More than anything, it's Intel's advantage in chip manufacturing that has defined the company. But only now is it ramping up production of 10nm silicon as part of its new 10th generation of Core processors.

Of course, this is Intel, and nothing is ever simple. Strictly speaking, Intel trickled out a few dual-core 10nm chips in 2018. But if you can point us at an available retail product that uses these ultra-rare Cannon Lake chips, you'd be smarter than us. 10th-generation Core CPUs thus represent the first true availability of 10nm processors from Intel. Likewise, while Intel's new 10th-gen line-up does include

both 10nm silicon and an exciting new microarchitecture, it also encompasses not just one but two 14nm designs based on existing microarchitectures. Intel 10th-gen, in other words, means a lot of different things.

That's just 10th-gen as we know it today. Don't be surprised if it becomes even more complicated. And that's just the products themselves. When it comes to branding, something Intel seemingly sees primarily as an opportunity to confuse its customers, the company is arguably achieving new lows with 10th-gen. The new nomenclature is so complicated and contradictory, it feels like you need a couple of MBAs, a law degree, and a PhD just to grasp the basics. But don't despair. There is some genuinely new technology buried beneath the brain-dead marketing. What's more, we're here to do the donkey work for you, and make sense of Intel's marketing madness. Let's begin.

**BEFORE WE DEEP-DIVE** into all those new 10th-gen chips, there's an elephant that needs escorting from the path ahead. We speak of the extraordinarily long delay to Intel's 10nm production process, the successor to 14nm in all its many forms. According to Moore's Law, 10nm tech ought to have been in full flow by late 2016. After all, Moore's Law famously predicts the doubling of transistor density every couple of years, and thus essentially dictates a new process technology from Intel every two years.

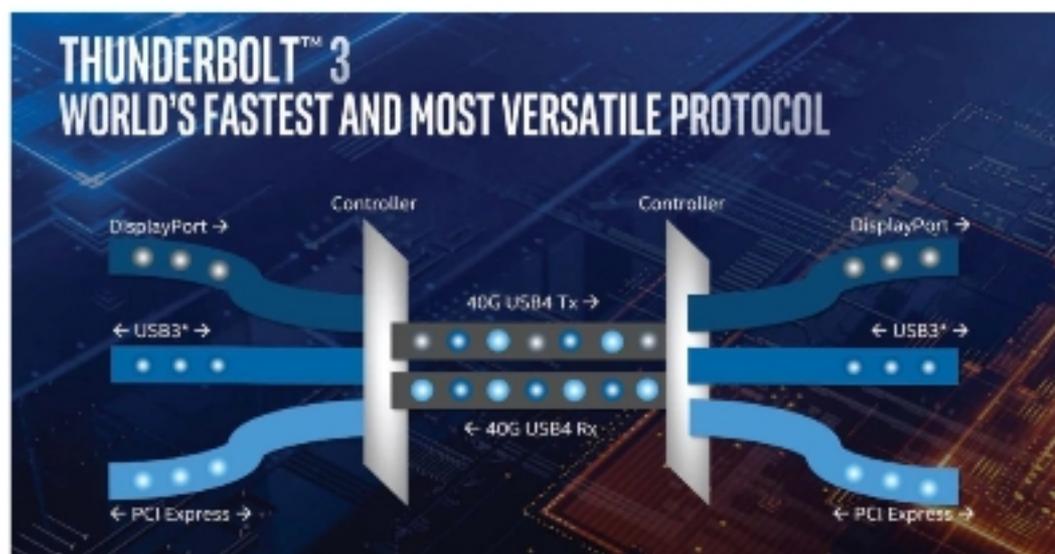
2016 sailed by, as did 2017. In 2018, Intel released a very small number of dual-core processors based on its 10nm node. But they were vanishingly rare and not true retail products. Even Intel described those initial 10nm products as "low volume." If we were cynics, we'd say they existed only so that Intel could claim that 10nm was in production and on sale.

Whatever, it's only now that 10nm CPUs from Intel are truly available, albeit with restrictions. At launch, 10nm is purely for mobile platforms. So, what happened? Inevitably, Intel hasn't been expansive in its explanations, but two fundamental issues have very likely been at play.

One: Intel was particularly aggressive with 10nm with regard to transistor density. For some time, the nomenclature attached to production nodes has become disconnected from the size of features in chips. In other words, the features inside a "14nm" Intel processor—the transistor gates, the logic cells, and so on—are not actually 14nm. By Intel's own figures, 14nm feature size actually ranges from 42nm through 399nm. That's the same for all major makers of integrated circuits. The node names have become labels, indicators of a given generation, rather than directly reflecting feature size.

Anyway, with feature size decoupled from the node nomenclature, the label "10nm" doesn't dictate transistor density. That's down to decisions made by Intel. For 10nm, Intel aimed for "hyper scaling" of density, in this case a 2.7x improvement in the number of transistors by area of integrated circuit. That compares with a scaling of 2.1x with the transition from 32nm to 22nm, and 2.5x improvement when 14nm came on stream.

Making matters worse, Intel went after this hyper scaling just as it was also running up against the limitations of its existing approach to producing chips using lithography based on ultraviolet (UV) light. It's the wavelength of the light used that informs the minimum size of features in a chip. UV light wavelength is in the order of 200nm. Using multiple masks and other tools, it's possible to produce feature



Integrated Thunderbolt 3 is part of the 10nm proposition.

10th-gen Core previews graphics technology that Intel will use to take on Nvidia and AMD.

sizes smaller than the actual wavelength. But with 10nm, Intel was pushing up hard against the limitations of UV light. 10nm is the end of the line for UV lithography. At the same time, the technology for extreme UV or EUV lithography, with wavelengths below 15nm, wasn't ready.

#### GENERATIONAL SHIFT

If that's the 10nm story, what about the 10th-gen narrative? While Intel's 10th-gen Core processors include new 10nm chips, there's much more besides. But we'll start with the 10nm hardware, because that's where the newest tech is found.

10th-gen chips that are 10nm also debut a new CPU architecture, known as Ice Lake. Depending on how you measure these things, Ice Lake represents the first extensively revised Intel x86 architecture since at least Skylake from 2015, and just possibly the first Intel Core chips in 2006.

For Ice Lake, Intel has cooked up both a new microarchitecture for its CPU cores, Sunny Cove, and new graphics tech, denoted Gen11 Graphics. Together they form the basis of

Intel's key CPU and graphics technologies going forward. The latter half of the equation includes Intel's upcoming discrete graphics cards, which are likely to be based on the same architecture as the Gen11 integrated GPU in Ice Lake.

Headline Ice Lake features start with a claimed 18 percent improvement in per core, per clock CPU performance compared to Skylake. Next up is a 2.5x increase in what Intel calls AI performance, the addition of AVX512 instructions, two times faster HEVC video encode, integrated Thunderbolt 3, DisplayPort 1.4a, and a whole lot more. In other words, Ice Lake is no simple die shrink with a few tweaks. It's the first

major overhaul from Intel in a long time, worthy of the 10th generation tagline.

For now, Ice Lake is limited to mobile applications and is exclusively available in dual and quad-core variants for low and ultra-low power applications. Thus, Ice Lake is targeted at thin-and-light laptops and below, the latter encompassing even slimmer form factors, such as tablets.

Arguably less worthy of the 10th-gen designation are the other members of the line-up. For now, the most straightforward 10th-gen CPUs are the new HEDT (high-end desktop chips), the beasts based on Intel's workstation and server technology. Known as Cascade Lake-X, they're 14nm chips, ranging from 10 to 18 cores.

Feature-wise, they're dead ringers for their Skylake-X predecessors, which offered essentially the same microarchitecture and the same core counts. However, the big news with Cascade Lake-X is pricing. While the 10-core Skylake-X chip was priced at \$989, the equivalent Cascade Lake-X model is just \$590. Similarly, Intel has slashed the price of the 18-core monster from \$1,979 in Skylake format to just \$979 for Cascade Lake-X. In short, Intel has pretty much cut pricing in half. The reason? AMD and the massively increased competition it is now presenting, especially in the form of its latest Ryzen 3000 series processors and imminent third-gen Threadripper CPUs.

The final part of the 10th-gen puzzle, for now at least, is another set of mobile CPUs, this time based on existing 14nm production technology, an architecture derived from Skylake, and designated Comet Lake. Just like the Ice Lake chips we mentioned a moment ago, these Comet Lake processors are low and ultra-low power models, but along with dual and quad-core variants, Comet Lake also marks the first time that Intel has offered low-power six-core configurations.

Yes, you read that right. You can now buy a laptop based on a brand new 10th-gen Intel Core processor. And if it's dual or quad-core, it may be low or ultra-low power, it might be 14nm or 10nm, it might be based on Skylake from 2015, or it might be a brand new architecture. Exactly how it makes sense that both of these families of CPUs, plus the HEDT multicore chips, are all designated 10th-gen is beyond our comprehension. But what we can say for sure is that it makes for a massive headache when it comes to actually recognizing these CPUs in shipping products, particularly with regard to the mobile CPUs.

Allow us to elaborate. Traditionally, Intel has designated its low and

## DARE WE MENTION DESKTOPS?

Intel's 10th-gen CPUs are here. But wait—what about the desktop? For now, there are three different families of 10th-gen chips and none are mainstream desktop processors. So, what gives?

According to the narrative that has emerged over the last year or so, yup, the problem is 10nm. Rumors suggest that Intel is still having problems clocking the new process up, and also with yields. The first issue seems to be confirmed by the low clock speeds and relatively unimpressive TDPs of the first 10nm Ice Lake mobile chips.

By way of example, the quickest quad-core low-power Ice Lake CPU tops out at 4.1GHz, while the fastest of the new 14nm

Comet family, also part of 10th-gen, hits 4.9GHz. In the ultra-low power area, Intel rates Ice Lake at 9W to Comet Lake's 7W. In some regards, those comparisons may be misleading. It's true that Ice Lake's Sunny Cove CPU cores bring more performance per clock, but the overall picture doesn't seem to be of a production process in rude health and delivering the typical benefits of a new node.

In a desktop context, of course, it's the clock speeds that are most worrisome. If the 10nm node really is the better part of a whole gigahertz behind 14nm in that regard, the 18 percent bump in per clock performance that comes with those Sunny Cove

cores will be neutralized by the lower frequencies.

What with the well documented overall delay to 10nm—around three years—it's not surprising that rumors based on supposedly leaked roadmaps and suggesting Intel may skip 10nm for the desktop and go straight to 7nm have emerged. However, Intel has explicitly debunked such stories. A recent official statement from Intel said, "We continue to make great progress on 10nm, and our current roadmap of 10nm products includes desktop." Of course, "current roadmap" isn't quite the same as "definitely being launched." We'll believe 10nm mainstream desktop chips when we see them.



Ice Lake is great for laptops, but what about desktops?

ultra-low power chips with "U" and "Y" signifiers, respectively. By way of example, the eighth-generation Core i5-8265U was a quad-core low power CPU, with a 15W TDP, and a maximum boost clock of 3.9GHz, while the Core i7-8500Y was an ultra-low power dual-core model, with a 4.2GHz boost clock, and a 5W TDP. Is it a little confusing to have a dual-core Core i7 and a quad-core Core i5? Perhaps. But the U and Y designations at least made it clear what class of CPU you were dealing with, so there was logic of a kind.

But not with 10th-gen. Oh, no. In total, there are currently 19 different 10th-gen

mobile CPUs to choose from. If we begin with the 10nm Ice Lake chips, the new nomenclature goes like this. Take the new Core i7-1065G7 as an example. It's a quad-core 10nm model with the most powerful "Iris" graphics option. Initially, the branding is fairly straightforward. The "i7" is Intel's established brand modifier indicating a high-performance model within a given range of CPUs. Next up, the "10" indicates a 10th-gen Core processor. So far, so good.

The next two digits, the "65," identify both the specific SKU in question and its status as a low or ultra-low power

## THOSE NEW 10TH-GEN CHIPS IN FULL

In what is increasingly becoming an Intel tradition, the new 10th-generation Core processor family is nothing if not utterly baffling. So, here's every one of the new chips, along with their key specifications.

### INTEL 10NM ICE LAKE-U SERIES MOBILE

	Cores/Threads	Base/Boost (GHz)	L3 Cache	TDP	GPU Units	GPU Clock (MHz)
Core i7-1068G7	4/8	2.3/3.6	8MB	28W	64	1,100
Core i7-1065G7	4/8	1.3/3.5	8MB	15W	64	1,100
Core i5-1035G7	4/8	1.2/3.3	6MB	15W	64	1,050
Core i5-1035G4	4/8	1.1/3.3	6MB	15W	48	1,050
Core i5-1035G1	4/8	1.0/3.3	6MB	15W	32	1,050
Core i3-1005G1	2/4	1.2/3.4	4MB	15W	32	950

### INTEL 10NM ICE LAKE-Y SERIES MOBILE

	Cores/Threads	Base/Boost (GHz)	L3 Cache	TDP	GPU Units	GPU Clock (MHz)
Core i7-1060G7	4/8	1.0/3.4	8MB	9W	64	1,100
Core i5-1030G7	4/8	0.8/3.2	6MB	9W	64	1,050
Core i5-1030G4	4/8	0.7/3.2	6MB	9W	48	1,050
Core i3-1000G4	2/4	1.1/3.2	4MB	9W	48	900
Core i3-1000G1	2/4	1.1/3.2	4MB	9W	32	900

### INTEL 14NM COMET LAKE-U SERIES MOBILE

	Cores/threads	Base/Boost (GHz)	L3 Cache	TDP	GPU Units	GPU Clock (MHz)
Core i7-10710U	6/12	1.1/3.9	12MB	15W	24	1,150
Core i7-10510U	4/8	1.8/4.3	8MB	15W	24	1,150
Core i5-10210U	4/8	1.6/3.9	6MB	15W	24	1,100
Core i3-10110U	2/4	2.1/3.7	4MB	15W	24	1,000

### INTEL 14NM COMET LAKE-Y SERIES MOBILE

	Cores/Threads	Base/Boost (GHz)	L3 Cache	TDP	GPU Units	GPU Clock (MHz)
Core i7-10510Y	4/8	1.2/3.2	8MB	7W	24	1,150
Core i5-10310Y	4/8	1.1/2.8	6MB	7W	24	1,050
Core i5-10210Y	4/8	1.0/2.7	6MB	7W	24	1,050
Core i3-10110Y	2/4	1.0/3.7	4MB	7W	24	1,000

### INTEL 14NM CASCADE LAKE-X SERIES HIGH-END DESKTOP

	Cores/Threads	Base/Boost (GHz)	L3 Cache	TDP	GPU Units	GPU Clock (MHz)
Core i9-10980XE	18/36	3.0/4.8	24.75MB	165W	N/A	N/A
Core i9-10940X	14/28	3.3/4.8	19.25MB	165W	N/A	N/A
Core i9-10920X	12/24	3.5/4.8	19.25MB	165W	N/A	N/A
Core i3-10900X	10/20	3.7/4.7	19.25MB	165W	N/A	N/A

processor, "5" in this case indicating low power—alternatively, "0" would signify an ultra-low power chip. Keeping up? Finally, the "G7" reveals the level of graphics power, with "7" designating the top tier of Gen11 graphics, with fully 64 execution units. If you've got all that, you'll no doubt immediately grasp that the Core

i5-1030G4 is a quad-core 10nm ultra-low power 9W Ice Lake processor with a top Turbo speed of 3.5GHz, and Gen11 graphics, with 48 execution units, right?

Now, while this is a break from the past, it follows a certain logic. So, all Intel needs to do is stick with that logic and we all have half a

chance of keeping up. Except that isn't what happens. For starters, even within the Ice Lake series, there's a Core i7-1068G7, which breaks the "5" and "0" indicators for low and ultra-low power. But that's nothing compared to the confusion that occurs when you add the alternate Comet Lake 10th-gen mobile processors into the mix.

That's because Comet Lake retains the old "U" and "Y" designators. Thus the Core i7-10710U is a low power 15W 10th-gen 14nm processor with six cores, Gen9.5 graphics, and a top Turbo speed of 4.7GHz, while a Core i5-10310Y is an ultra-low power 7W quad-core chip, again with Gen9.5 graphics, and a maximum Turbo speed of 4.1GHz.

Quite how Intel imagines mainstream consumers can follow any of this, we have no idea. It's not entirely clear, for that matter, why Intel felt the need to offer two parallel low and ultra-low power CPU series within the 10th-generation line-up. However, we suspect that concerns over 10nm production capacity are at the heart of the matter.

The next obvious question is which of the two 10th-gen mobile families should you go for. That is a very tough question to answer, unless you either demand six cores or want the best possible integrated graphics performance. In the first scenario, only 14nm Comet Lake can deliver. In the second, you want an Ice Lake Processor with its fancy new Gen11 graphics. But for everything else? It's unclear at best.

On the one hand, Ice Lake offers a more efficient architecture that does more work per clock cycle. On the other, Comet Lake SKUs tend to clock higher. As for battery life, that's confusing, too. You might expect 10nm tech to be clearly superior for battery life, but the 10nm Ice Lake ultra-low power models, for instance, are rated at 9W, while the equivalent Comet Lake options are pegged at 7W.

In practice, your choice will often be dictated by other features, such as screen technology, chassis quality, and so on. But if you've narrowed down your options and are left with both flavors of 10th-gen CPU and little else to base a decision on, we'd lean toward the 10nm option for its stronger feature set. Extras such as integrated Thunderbolt 3 and support for DisplayPort 1.4a maximize connectivity options, particularly when it comes to adding future high-res displays.

Back on the desktop, of course, it's a lot easier. On the HEDT side, the pricing of 10th-gen Cascade Lake-X makes it a massive step

## GET READY FOR SOME DECENT GRAPHICS

Over one teraflop of pure compute performance, support for ultra-high resolution HEVC encode and decode, improved display pipelines, an enhanced rasterizer, variable rate shading, over double the performance in 3D rendering, and support for adaptive sync. This, ladies and gentlemen, is the elevator pitch for Intel's new Gen11 graphics technology, as seen in the new 10nm Ice Lake mobile CPUs.

In terms of actual in-game frame rates, Intel reckons that the fastest Gen11 integrated graphics core is typically around 80 percent faster than the Gen9.5 graphics found in eighth-generation processors, and sometimes over 100 percent faster. Much of

that improvement is down to increased complexity. Gen11 graphics tops out at 64 execution units to the maximum 48 of its Gen9.5 progenitor. However, the 48-execution-unit iteration of Gen9.5 was limited to a small number of SKUs, whereas the 64-unit Gen11 graphics is found in nearly half of the currently available Ice Lake processors, including ultra-low power versions.

3D performance aside, Gen11 brings some other very nice enhancements. It supports two HEVC 10-bit video encode pipelines in parallel, and in terms of display output, Gen11 can drive three 4K screens at the same time—two via DisplayPort 1.4 HBR3, and another courtesy of HDMI 2.0b. Alternatively, it can drive dual 5K displays or a

single 4K display at 120Hz. Nice. Very nice.

But what is the relationship between Gen11 and Intel's new discrete graphics chips, due next year and known as Intel Xe? Officially, we don't know. But our expectation is that Xe will be based on the same building blocks as Gen11, but with some additional enhancements, including hardware support for ray tracing. For those discrete cards, the indications are that initial SKUs will include GPUs with 128, 256, and 512 execution units. Given that Gen11 with 64 units is broadly good for 1080p gaming at 30fps and decent detail levels, the higher-end Xe chips, with nigh on 10 times the complexity, make for an interesting proposition.



Is 10th-gen Core where Intel finally gets in on the game?

forward in terms of value. You get double the cores at any price point compared to Skylake-X. As for mainstream desktop, there are no 10th-gen chips at all, for now.

Overall, then, this is the most interesting new generation of chips from Intel in a long time. But also the messiest. It includes Intel's first new production node in five years, and its most significant architectural overhaul in at least four years, not to mention debuting what will probably be the basis of Intel's

graphical assault on Nvidia and AMD. But amalgamating the genuinely new 10nm products with a bunch of rebaked 14nm chips under a single 10th-gen banner, then cooking up some of the most baffling branding yet devised, is also a source of frustration. At this point, the term "generation" as Intel uses it means pretty much zilch. Still, you now know exactly what to expect from the CPUs Intel has deemed to describe as 10th-generation Core processors. For a while, at least. ☺

# SUBSCRIBE TO **MAXIMUM PC**

THE GO-TO RESOURCE FOR PC ENTHUSIASTS...

**CHOOSE YOUR PERFECT PACKAGE**



**1**

THE DIGITAL EDITION  
**ONLY \$9.00**  
EVERY YEAR

Instant digital access on your iPad, iPhone, and Android device.



**2**

THE PRINT EDITION  
**ONLY \$24.00**  
EVERY YEAR

Every issue delivered to your door for a fraction of the store price.

3

COMPLETE PRINT  
& DIGITAL BUNDLE  
**ONLY \$33.00**  
EVERY YEAR

Get the print edition of *Maximum PC* delivered direct to your door, and instant access on your iPad, iPhone, or Android device.



THE EASY WAYS TO SUBSCRIBE...

<http://bit.ly/sub-maxpc>

Or call 844-779-2822 (toll-free)

**TERMS AND CONDITIONS** Prices and savings quoted are compared to buying full-priced US print and digital issues. You will receive 13 issues in a year. You can write to us or call us to cancel your subscription within 14 days of purchase. Your subscription is for the minimum term specified, and will expire at the end of the current term. Payment is non-refundable after the 14-day cancellation period, unless exceptional circumstances apply. Your statutory rights are not affected. Prices correct at time of print and subject to change. Call center opening hours are Monday to Friday 8am-7pm, and Saturday 10am-2pm EDT. For full terms and conditions, please visit <http://bit.ly/maxpc>. Offer ends: January 31, 2020.





# TAKE CONTROL OF YOUR AUDIOBOOKS



Rip, convert, and stream your audiobook collection to any device, with *Nick Peers*

**WHEN IT COMES TO MEDIA**, your PC has everything covered, right? There's Kodi, Plex, or Emby for starters, managing your music, movies, and TV with aplomb. Except there's something missing....

Everyone loves a good story, and audiobooks (along with radio-style dramatizations) provide an easy way to devour books when focused on other tasks, whether driving long distances or tweaking your PC's hardware. The obvious place to go for an audiobook fix is Audible, but it ties you into a proprietary ecosystem that requires you to purchase your books outright via one-off fees or a recurring subscription. Furthermore, you're indebted to Audible's own apps and servers for accessing your

books. That's fine if you're a fan of paying for streaming services, but the likes of Plex, along with audio-only services such as Logitech Media Server, enable you to stay in complete control of your media, so why shouldn't that be the case with audiobooks? There's no reason. You can incorporate your audiobook collection into existing streaming services, but none of them quite hits the mark. Which is where we come in.

In this exhaustive guide, we'll start by bringing all your audiobooks—whether CD, free downloads, or purchased through Audible—under one roof. You'll rip CDs exactly the way you want your books divided up, then tag your files consistently (complete with high-quality embedded

artwork) to make browsing, managing, and listening to audiobooks a pleasure.

Once your media is in place, we'll reveal how to set up a dedicated audiobook server on your PC or NAS, enabling you to stream your audiobooks to any device in your home—or further afield if you want access on the go over the Internet. There's even offline access through your cell phone should you wish to be able to listen to your books in the most remote spots.

We've spent months developing this project, experimenting, tweaking, and tearing our hair out so you don't have to. It's time to give your audiobook collection the prominence it deserves, so simply turn the page to put our carefully crafted research to good use.



**YOUR FIRST STEP** is to bring all your audiobooks under one digital roof. Start by creating a single folder—"Audiobooks"—somewhere accessible, such as inside your Music folder. Inside here you'll store your audiobooks using the following folder structure: Collection\Author\Book. In other words, each book has its own folder, inside which your audiobook can be housed in multiple files (such as individual chapters, episodes, or parts), to aid navigation and playback.

As an aside, the Collection folder level is optional, designed to work with your audiobook server's support for multiple users, so you can restrict access to books on a per-user level. More on that later.

Next, you want to migrate all your audiobooks into this folder structure. If you're an Audible user (past or present), check out the box below on using OpenAudible to download DRM-free MP3 copies of your purchased books to incorporate into your new setup.

You can also obtain audiobooks from other online sources—books in the public domain (typically out-of-copyright classics) have been narrated by volunteers and placed on various websites, many of which have been produced by LibriVox, so head to <https://librivox.org> to browse its catalog. Choose the "Whole book [zip file]" option to download the book as a

series of MP3 files, one per chapter. Each file is tagged with author, title, and other key details, and you can also download cover art to embed into the files. More on verifying those tags and embedding the artwork later. For now, once downloaded, extract the contents of the zip file into your Audiobooks folder according to your folder structure: for example, Audiobooks\Classics\Herman Melville\Moby Dick.

### RIP YOUR AUDIOBOOK CDS

Another way to add audiobooks to your library is to rip them from CD. We recommend iTunes (install it through the Microsoft Store) for three reasons: It has a built-in audiobook library to avoid cross-contamination with your music; its CD lookup database includes many popular audiobooks, from *Harry Potter* to Stephen King; and it can join multiple tracks into a single file when ripping. This helps mitigate those audiobook CDs split into very short tracks to make them easier to navigate on CD players.

After installing, open iTunes, and insert your first audiobook CD. After a short pause, you'll hopefully see the "CD Lookup Results" dialog appear, with at least one match to choose from. Select the first in the list, and click "OK." Review how it's named each track—ideally, you'll see the chapters clearly marked, but if not, click

the "Options" button (the cog next to the eject button), and choose "Get Track Names" to choose another from the list.

If you can't tell which tracks correspond to specific chapters, double-click a track name to start playing it, to see if the chapters are spoken at the start of a track. Once identified, select all the tracks that make up chapter one, click "Options," and select "Join CD tracks"—you'll see the selected tracks are now all linked to indicate they'll be ripped as a single file. Repeat for other chapters on this disc.

Before ripping, click "CD Info" to review the audiobook's key metadata—you'll also need to click this to manually enter the information if the CD lookup database found no matches. Verify the "Artist" is the author, and if necessary, tweak the audiobook title in the album field—or "album" name if it refers to multiple discs or parts. This needs to be identical across all the audiobook's discs to ensure they're ripped to the same folder. Use the "disc number" fields to specify the disc number. If the genre is currently set to "Spoken & Audio," change this to "Audiobook," unless it's a dramatization rather than a straight reading. Finally, verify the compilation box is unchecked before clicking "OK" followed by "Import CD."

The "Import Settings" box pops up. Select "MP3 Encoder" for the "Import

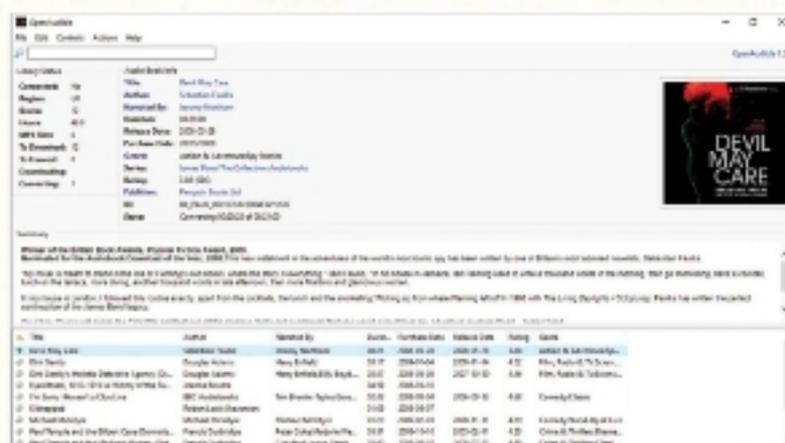
## DOWNLOAD AUDIBLE BOOKS

To add your Audible collection to your new DRM-free audiobook collection, you need OpenAudible, a tool for managing and downloading Audible books you own to your PC. Head over to <https://openaudible.org> to download the setup file. Launch it, clicking "More info → Run anyway" when the SmartScreen filter pops up. Once installed, launch OpenAudible, and press Ctrl-K to connect to your Audible account using OpenAudible's built-in browser. Close the window and, if necessary, press Ctrl-K to connect again, then choose "Controls → Full Library Sync" to view all your

previous purchases in the main OpenAudible window.

By default, all downloaded audiobooks are converted to MP3 (verify this via "Edit → Preferences"). To download a single audiobook, right-click it in the lower pane, and choose "Download." Keep an eye on the "Audio Book Info" pane, which tracks the progress of the download and conversion.

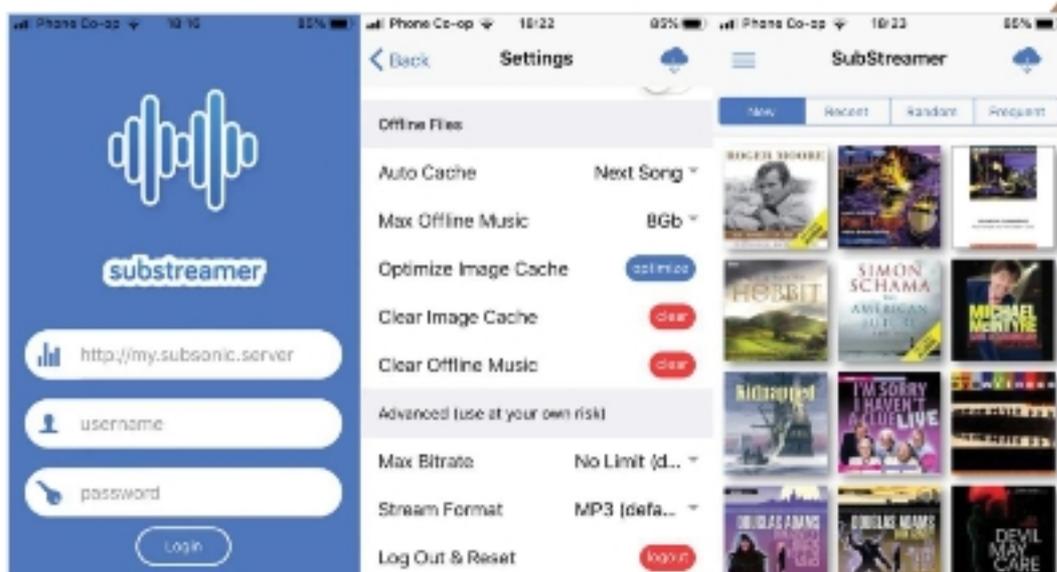
OpenAudible downloads audiobooks as single files—if you want to divide them into parts, choose "File → Split by Chapter." Browse to C:\Users\{username}\OpenAudible\mp3 (substitute {username} with your username), where you'll find all the converted files. Select one and click



"Open," then select your main Audiobooks folder, and create the required Author/Title subfolders before selecting the Title and clicking "OK." If Audible has chapter data, your files are split accordingly.

Once complete, switch to MP3Tag, and follow the advice in the main text to renumber the tracks correctly, then automatically retitle and

rename each file according to how you want to do so (such as "Chapter 01" or "Part 02"). Finally, the artwork has only been embedded into the first title, so select title 1, right-click the artwork, and choose "Copy cover." Now select the remaining titles, right-click the artwork again, but this time choose "Paste cover" before clicking "Save."



**Audio-Technica ATH-MSR7—high-res audio headphones from the aural heavyweight.**

The SubStreamer app for iPhone works beautifully with Booksonic servers.

Using" box, and click "Custom" under "Setting." Audiobooks don't require top-quality settings, so minimize the space they take up by setting the "Stereo Bit Rate" to 64kb/s. Leave the other settings as they are and click "OK" twice. Wait for the disc to be ripped, then repeat for any other discs in the audiobook set.

#### JOIN CHAPTERS

If you're ripping a multi-disc audiobook, there's a good chance some chapters will be split across two discs. You can either leave them as they are (say Chapter 10a

and Chapter 10b) or use a tool such as Free MP3 Joiner ([www.convertaudiofree.com/free-mp3-joiner.html](http://www.convertaudiofree.com/free-mp3-joiner.html)) to combine them into a single file. Free MP3 Joiner is a tiny portable tool that doesn't reconvert the files—it just stitches them together. Launch the app, then click the first "Browse" button to select the file containing the first half of the chapter before clicking the second "Browse" button to select the second half. Finally, choose a suitable output name (such as "Chapter 10"), and click "Join." Once complete, click "Open Output folder" to

verify the new file exists alongside the original two, which you can now delete.

By default, iTunes rips all CDs to its own iTunes\iTunes Media\Music folder. Navigate here to find a folder named after your audiobook's author. Open this to reveal the audiobook folder, which you should move to the Audiobooks\ Collection\Author folder you made earlier.

#### TAG YOUR AUDIOBOOKS

Once you've ripped your first audiobook, you need to tidy up the tags. Because audiobooks use the same tagging

© SUBSTREAMER, BOOKSONIC

## BUILD A LOCAL AUDIOBOOK LIBRARY

Having brought all your audiobooks together, now all you need to do is play them. If you only plan to listen to them on your PC, both iTunes and MusicBee (<https://getmusicbee.com>) offer dedicated audiobook views, enabling you to not just split your audiobooks from your music files, but also take advantage of specific audiobook features, such as the ability to pick up from where you left off during playback.

If you plan to go down the iTunes route, we recommend moving the Audiobooks folder back inside the default iTunes Media\Music folder. After being added to your iTunes

library, you'll find them automatically classified as music. You can move them across by right-clicking a title (or titles) and choosing "Get Info → Options tab." From here, change its "Kind" to "Audiobook." It then vanishes from the Music library (switch to Audiobooks view to get it back). Sadly, the way iTunes works is to sort the audiobook library by title, so multiple entries for each audiobook will appear. To fix this, download the AudiobookSortName script from <http://samsoft.org.uk/iTunes/scripts.asp>—once done, select all your audiobooks in iTunes, then double-click the script. Click "Yes," and wait for a minute



or so for it to perform the sort silently in the background. If your security software protests, you should either allow the script to run, or temporarily create an exception for C:\Windows\System32\wscript.exe.

MusicBee also has a dedicated audiobooks view—after adding your audiobook folder to the library, you

need to right-click each author or narrator in turn, and then choose "Send to → Audiobooks" to transfer them to MusicBee's Audiobooks view. The primary advantage of doing this is that MusicBee will remember where you stopped listening to an audiobook, and then pick up from where you left off the next time you open it.



system as your music files, you have to “translate” author and narrator information to the correct music-based tags. First, you need a tagging tool—our go-to app is MP3Tag, which you can download and install (or run as a portable app) from [www.mp3tag.de/en](http://www.mp3tag.de/en).

Launch MP3Tag and browse to `Audiobook\Collection\Author`. Highlight your audiobook folder and click “Select Folder” to open it. All your ripped files should appear in the main pane. Start by reviewing the universal tags in the left-hand pane: Select all the tracks and ignore “Title” (which says “<keep>,” indicating it contains multiple variables for each individual file). Instead, focus on “Artist,” “Album,” “Album Artist,” and “Composer.” We recommended putting the narrator’s name in the “Artist” field, and the author’s name in the “Album Artist” and “Composer” fields. Verify the “Album” is the book’s title. Make any amendments, change the “Discnumber” field to “<blank>,” and click the floppy disc icon in the top-left to save your changes.

You can now turn to the files that make up your audiobook. Start with the “Track” field in the main pane (scroll right to access it)—this needs to be renumbered sequentially from top to bottom, and the quickest way to do this is to select all your tracks, then choose “Tools → Auto-numbering wizard.” Make sure “Tracknumber” is checked and it’s set to begin at track number 1, then click “OK.”

Now review the “Title” field. Each title is an amalgamation of the titles of

tracks you’ve joined. To change this manually, click inside the first title, then type your replacement title, and press Enter to move on to the next title. If, however, you simply want to name each title for its component chapter (Title—Chapter 01, Title—Chapter 02, and so on), speed things up by creating a custom action.

Select all your tracks and choose “Actions → Actions.” Click “New,” name your action group “Quick chapter titles” or something similar, and click “OK.” Click “New” again, and choose “Format value” as your action type, followed by “OK.” Select “TITLE” for the field, and type the following into the “Format string” box before clicking “OK”:  
`%album% - Chapter $num(%track%,2)`

Click “OK” again to return to the “Action groups” window, check “Quick chapter titles,” and click “OK.” You’ll see each title renamed accordingly: Book title—Chapter 01, Book title—Chapter 02, and so on. Don’t want the leading zero? Type `%title% - Chapter %track%` into the “Format string” box instead. The action can be duplicated and modified to replace “Chapter” with, say, “Episode” or “Part.”

Finally, we want to rename the file names. We’re using the convention “The Colorado Kid—01.mp3,” “The Colorado Kid—02.mp3,” and so on. Make sure all the tracks are selected, then choose “Convert → Tag—Filename.” You’re prompted

**Booksonic has a built-in player, accessible through its web interface.**

to select a Format string—type the following into the box, and click “OK”:  
`%album% - $num(%track%,2)`

#### ADD ALBUM ART

Now your audiobook files are correctly tagged, there’s just one thing missing: artwork. To embed this into each file, select all your tracks, right-click the blank artwork box, and choose “Add cover.” Once selected, the artwork should appear in the box, at which point click “Save.”

We recommend adding a square image in JPEG format of the audiobook’s artwork that’s 600x600 pixels. Where can you get this? You could scan in your CD cover, or look online at <https://images.google.com>; type the title of the book along with “audiobook” and see what comes up. Click “Tools,” then “Size → Large” to filter out smaller images, then roll over each result to see what size image is available. Click one, then right-click it, and choose “View image” to bring it up in your browser. Save this to your Downloads folder, then open it in an image editor if necessary, to reduce its size to 600x600 before saving it in JPEG format, ready to embed into your file.

#### BUILD YOUR AUDIOBOOK LIBRARY

While it’s possible to keep your audiobook library on a single PC—see the box on building a local library, pg. 49—where’s the fun in that? Booksonic (<http://booksonic.org>) is an audiobook streaming server that enables you to stream audiobooks to any device on your network. It’s a fork of Subsonic, a well-known music streamer, so it can be used on a wide range of devices, including Android (Booksonic app, \$3), iOS (SubStreamer, free), and Windows (MusicBee—see box, pg. 49).

Your main choice is where to install the server: There are installers for Windows



# TECH PORN



## SHIELDS UP

Twenty-three USB ports. Twenty-three! Are you kidding us, MSI? Yes, the X570 Creation has a whopping 23 USB ports, 14 of which are situated on the rear I/O, which comes with its own pre-installed shield. One of those is a high-speed M.2 slot, and the remaining slots are supported via connectors on the board itself.



### 3 XTRA, XTRA!

This motherboard comes with an added bonus: an M.2 Xpander-Z card, which plugs into one of the board's PCIe slots to provide two more M.2 slots, complete with fan. The Xpander is a seriously well-made piece of hardware, with solid metal for heat spreading and a dedicated fan control switch.

### 2 ULTRA-COOL

The on-board fan uses a dual ball-bearing setup to produce minimal noise, even when any attached SSDs are running at load. The same goes for the fan on the M.2 Xpander card, so even if you've got both running, there's no need to worry. Between this and the three chunky heatsink modules, heat management is completely sorted.

# MSI Prestige X570 Creation

**HOW MUCH MONEY** should you spend on a motherboard? More, says MSI's Prestige series. No, more than that, too. If you're the sort of person who relishes the idea of dropping half a grand on your motherboard, this could be the board for you. It's part of the X570 selection of boards, and as such comes with all the latest bells and whistles, most importantly the introduction of the new PCIe 4.0 connections for M.2 drives and expansion cards.

We've also got Wi-Fi 6, 10G super LAN for incredible intra-network transfers, and some extremely sophisticated multiphase power design that provides maximum performance for overclocked CPUs. Heat management is perfectly handled, with steel guards for the DIMM slots, and a massive heat pipe that almost curves around the entire board. Then there's the main heat plate, which has a stylish metallic design (which matches the rear I/O shroud), and conceals two superfast Lightning M.2 slots with sticky heat pads.

Honestly, there are too many features for us to discuss here. We could talk about MSI's Zero Frozr technology, supporting this motherboard's integrated fan, or the support for AMD's Turbo USB feature, or even Audio Boost 4 functions. As the name implies, this board is perfectly suited to creators who demand a powerhouse as the beating heart of their system. The BIOS includes a "creator center," which enables users to switch between different system modes for optimized performance, whether you're a video editor or you work on modeling 3D architecture. —CHRISTIAN GUTTON

# ALL THE TUTORIALS, AND ADVICE YOU NEED TO SUCCESSFULLY INSTALL AND USE UBUNTU.

Linux is taking over the world, and the last frontier is your desktop. Find out what all the fuss is about with the ultimate guide to the most popular distribution, Ubuntu.



Ordering is easy. Go online at:

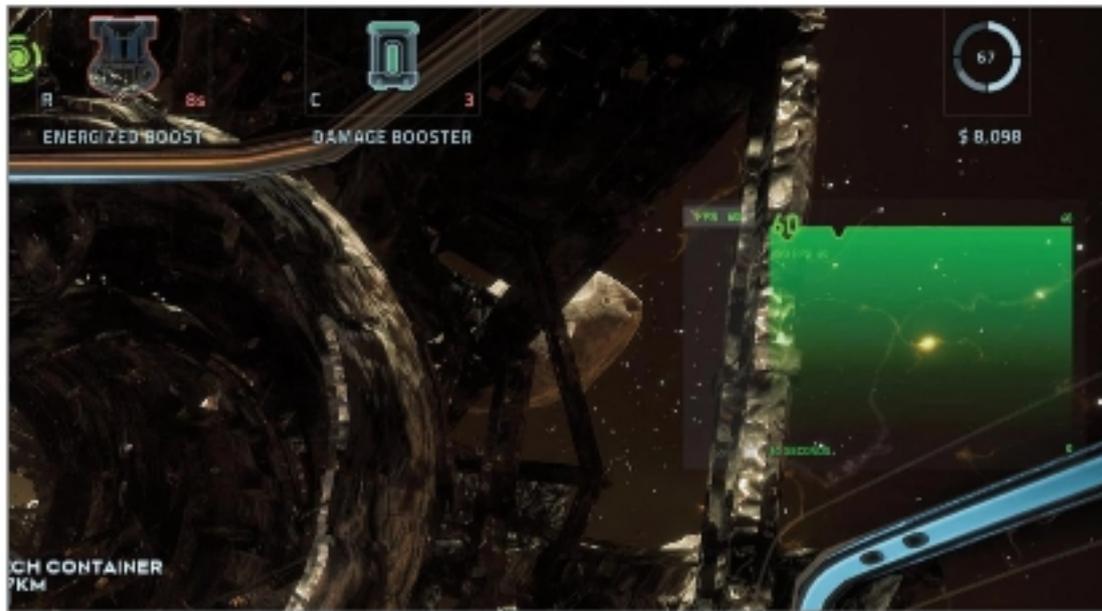
[www.myfavouritemagazines.com](http://www.myfavouritemagazines.com)

Or get it from selected supermarkets, newsagents, or bookstores

# HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

## TIP OF THE MONTH



### CHECK FRAME RATES WITH XBOX GAME BAR

Xbox Game Bar has had a lot of useful features added over the years, to the point where it's actually pretty good now. It's got achievement support, system metrics, and capture tools—all helpful for the serious gamer. The latest addition is a real-time FPS counter, which can be pinned as an overlay in just about any game. It charts a graph, too, demonstrating how stable your frame rate is.

## MAKE – USE – CREATE



**58**  
Advanced  
VirtualBox tips  
and tricks



**62**  
High refresh rate  
monitor tweaks for  
competitive gamers



**70**  
Enhance your  
video soundtracks  
with Audacity



CHRISTIAN GUYTON  
STAFF WRITER

### HEATSUNK

I've mentioned this before, but it aggravates me, so I'll do it again. What's with the current obsession with stuffing bulky heatsinks on to M.2 drives? The last drive I looked at had RGB lighting and wouldn't fit between the PCIe slots on my motherboard without removing half its plastic coating. Elsewhere, I see minuscule screws to remove a solid block of copper, or clips that appear easy to dislodge, but shred my fingertips.

Motherboard manufacturers are at it, too. Every new mid-range or better mobo seems to have its own flashy new cooling solution for M.2 drives, with dedicated fans on X570 motherboards, and every variety of thermal guard under the sun.

This is swiftly becoming a rant, and I'm OK with that. Why do some motherboards come with a single full plate that covers all the M.2 slots as well as the integrated fan? Worse still, if the manufacturers wanted to slap RGB lighting on top of that plate, there's going to be cables with tiny connectors to contend with as well. I'm sick of it. If I'm buying a high-end M.2 SSD, rest assured I can afford a high-end motherboard. On the other hand, more affordable boards and drives seem totally lacking in the heatsink department. We can't win—unless you know exactly who I have to bother at Samsung to get this nonsense sorted out....

submit your How To project idea to: [comments@maximumpc.com](mailto:comments@maximumpc.com)

GAME: REDDEAD REDEMPTION 2



presents:

# AUTTOPSY

THIS MONTH WE DISSECT...

## Fairphone 3



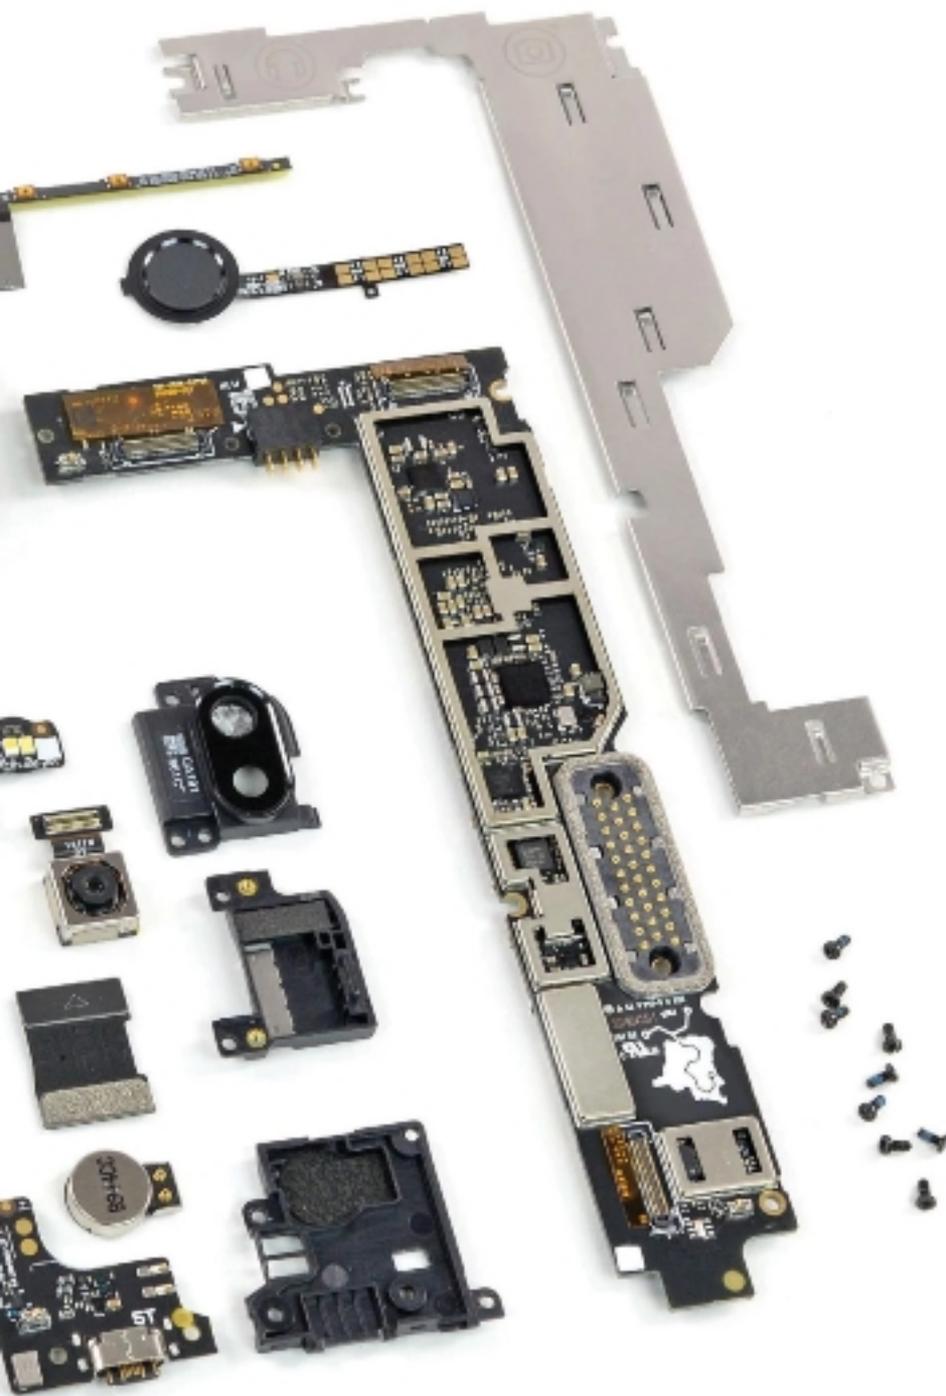
**About iFixit** iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their own products. To learn more, visit [www.ifixit.com](http://www.ifixit.com).



Removing the mobo, we discover a bunch of spring contacts.



The starting point for this month's teardown is a working Fairphone 3.



You can tear many of the modules apart as well—a real boon for fixing problems.

© FAIRPHONE

## BACKGROUND

"Good things come in threes," the Dutch social enterprise must have thought when planning the Fairphone 3, its newest ethically made mobile device, sporting a modular and repairable design.

## MAJOR TECH SPECS

- 5.65-inch IPS display with 1080x2160 resolution (427 ppi)
- Qualcomm Snapdragon 632 SoC
- 4GB of RAM and 64GB of storage
- 12MP f/1.8 rear camera and front-facing 8MP f/2.0 camera
- 3,060mAh battery
- Fingerprint sensor, headphone jack, and IP54 rating

## KEY FINDINGS

- We start by removing the back cover—no tools needed. Compared to the bumper-like casing of the Fairphone 2—as well as the two-piece "slim" version that replaced it midway through the product cycle—this is definitely an improvement. If you liked the added protection offered by the bumper, no worries—it's included in the box.
- Removing the battery is also easy as Android 9—it lifts out using the notch at the bottom. With recessed contacts and sturdier plastic casing, it's a well-protected throwback we're happy to see. At 11.781Wh (3,060mAh at 3.85V), this juice box is not the biggest around, but it should get you through the day—and if it doesn't, you can carry a spare and swap it out. It beats the Google Pixel 3a in capacity (11.55Wh) but can't quite touch the iPhone Xs Max (12.08Wh), let alone the Shift 6m (16.3Wh). The underside of the battery bears an encouraging message: "Well done. You're what progress looks like." Speaking of which, it's time to make some more teardown progress. Onward!
- Fairphone decided to forgo the nifty display lock mechanism from the previous iteration, opting for standard Phillips #00 screws instead. There's even a tiny screwdriver provided. With the screws gone, we unsnap the display from the frame—revealing the same pogo pin connector plus breakout board configuration seen in the Fairphone 2. Note the map of the Democratic Republic of Congo here, illustrating the conflict-free sourcing of minerals such as tantalum, tin, tungsten, and gold.
- The Fairphone 3 comes with a family of modules: top module, camera module, and bottom module, much like in the Fairphone 2. The loudspeaker now gets its own closed module (as opposed to being part of the bottom module), and connects to the rear of the bottom module. If you hope to upgrade your Fairphone 2 with these modules, it's bad news: Modules and their inner parts are not compatible. But their housings are said to be produced from 50 percent recycled polycarbonate, so there's some reuse.
- Repairability Score: 10 out of 10 (10 is easiest to repair). Key components like the battery and screen have been prioritized in the design and are accessible either without tools or with a regular Phillips screwdriver. Visual cues inside the phone help with disassembling and replacing its parts and modules. Replacing complete modules is very easy. Going for their internal parts is also possible and requires a Torx screwdriver. Replacement guides and spare parts are available via the manufacturer's website. Most components inside the modules are individually replaceable, although some are soldered on. ⚡

# Advanced VirtualBox Tips and Tricks

## YOU'LL NEED THIS

### VIRTUALBOX

Download from [www.virtualbox.org](http://www.virtualbox.org) for your Windows or Linux PC.

**EVERYBODY LOVES A VIRTUAL MACHINE**, don't they? Whether you're looking to try out an alternative operating system or want to resurrect an old version of Windows for compatibility purposes, VMs are the answer. And thanks to VirtualBox, the underlying technology powering that answer doesn't cost you a single cent.

VirtualBox is pretty straightforward to use, and by now you've probably got a handful of VMs up and running to cover your various bases. But perhaps performance isn't quite what you'd like it to be, or you're wondering whether there are other uses for virtual machines you've not thought to explore. In this tutorial, we're going to delve into some more advanced topics to make sure your virtual machine is performing at its best.

There's stuff to speed up your VM's performance, such as ensuring you've enabled hardware virtualization in the BIOS, plus tips and tricks for utilizing your hardware better, from setting up multiple displays to fine-tuning your VM's network settings to give it just the right amount of access it needs. We'll examine ways to resize your virtual disk images (as well as create images from physical drives) and introduce you to the command-line VBoxManage tool. Enough chit-chat, it's time to step your VirtualBox know-how up to the next level. —NICK PEERS



## 1 ENABLE VIRTUALIZATION

Most modern processors, whether Intel [VT-x] or AMD [AMD-V], support virtualization technology. This helps virtual machines tap directly into the hardware to help improve their performance, and only a fool would run VirtualBox without them. You can quickly verify that virtualization is both supported and enabled on your PC by right-clicking the Taskbar and choosing "Task Manager." Switch to the "Performance" tab (click "More details" if necessary), and check the "Virtualization" setting under "CPU." If it's absent, your processor doesn't support the tech (time for an upgrade?). If it's there, but listed as "Disabled," you need to consult your mobo manual for the setting you need to enable in the BIOS—for example, on the MSI X470 Gaming Plus mobo for AMD Ryzen CPUs, you need to go to "Overclocking → CPU Specifications," and enable "SVM Mode" [Image A].

» Once enabled, open VirtualBox, select a virtual machine, and click "Settings." Select the "System → Acceleration" tab, where you can check the "Enable VT-x/AMD-V" box. If you also see an "Enable Nested Paging" option, make sure it's also checked to allow your hardware to handle certain memory management functions, and offer an additional boost.

» There's also a "Paravirtualization interface" drop-down, set to "Default" by, er, default. This ensures that Hyper-V is selected for Windows guests and KVM for Linux guests. The "Legacy" option should be avoided in all cases except where your VM was created years ago in a version of VirtualBox before 5.0.

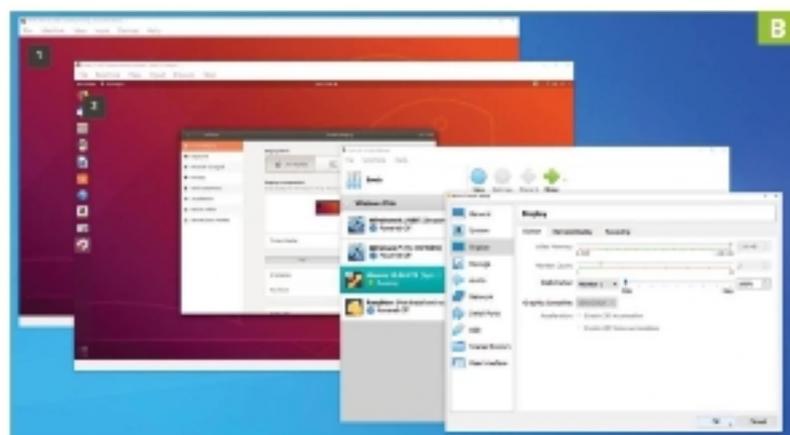
## 2 SET UP MULTIPLE DISPLAYS

Your virtual machines can—OS permitting—handle multiple displays, just like your main PC. These can be run as two virtual windows within a single display, but it's more practical to run them on two physical displays, as you would your regular desktop.

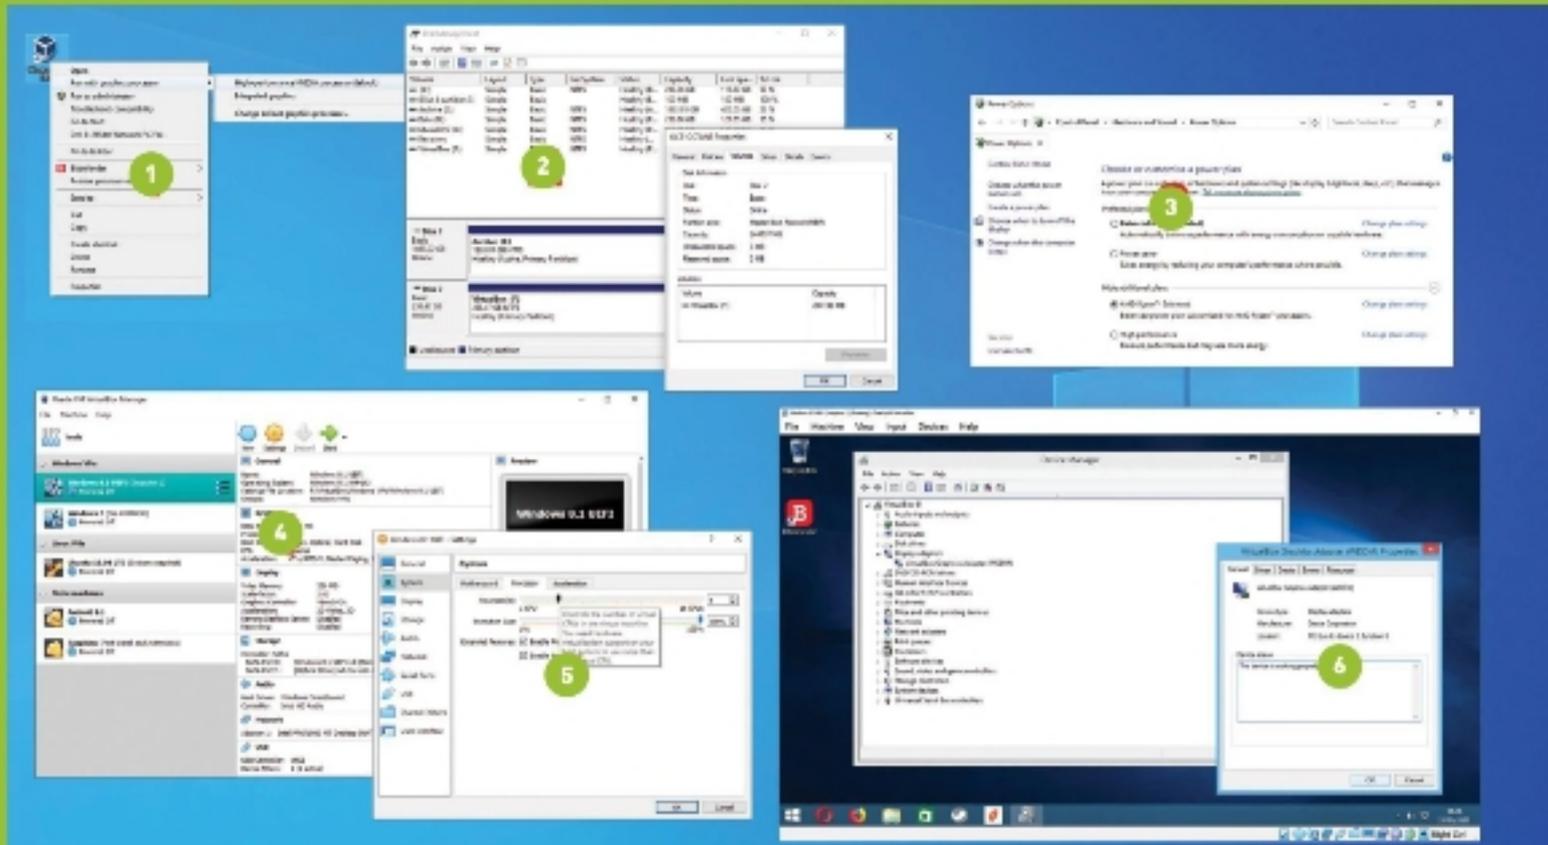
» Start by visiting the "Display → Screen" section of your VM's settings. Make sure the maximum RAM [128MB or 256MB depending on the guest] has been assigned to your VM's video memory, then push the "Monitor Count" slider to 2. There will only be a green bar between 1 and 2 if you have two monitors attached to your PC, both are switched on, and the desktop has been extended between them.

» You'll see a "Scale Factor" slider setting, too—this is primarily for higher-res displays. Ignore the recommended 175 percent setting for your 4K display; we find 200 percent is much clearer. You can configure the scale factor independently for differently sized displays: Just click the "All Monitors" drop-down to select each display in turn. Also verify the VBOXSVGA adapter has been selected as the graphics controller.

» Once configured, start your virtual machine and press Right Ctrl-F to run it full-screen. If it's showing on the wrong display, roll the mouse to the bottom-middle



# IMPROVE PERFORMANCE IN VIRTUALBOX



## 1. SWITCH GRAPHICS CARD

If you've got two graphics cards installed, right-click the VirtualBox desktop shortcut, and look for a "Run with graphics processor" option—make sure you choose your more powerful graphics chip.

## 2. RUN VMs ON AN SSD

Avoid trying to run virtual machines from platter-based hard disks—set up a dedicated SSD for all your VMs, and cut down on your thumb-twiddling time.

## 3. CHANGE YOUR POWER PLAN

Open "Settings → System → Power & sleep," and click "Additional power settings" to choose the "high performance" plan; you may find a CPU-specific option (such as "AMD Ryzen Balanced").

## 4. UPGRADE YOUR RAM

Now's the time to upgrade your PC to 16GB or even 32GB of RAM, allowing you to allocate a decent amount to your VM (or run multiple VMs at once).

## 5. ADD MORE CORES

Open your VM's "Settings → System → Processor" tab to allocate more processors to your VM—quad-core or even octa-core is possible with a CPU possessing 16 threads.

## 6. INSTALL GUEST ADDITIONS

Make sure the latest version of the Guest Additions is installed, as these install graphics drivers that enable you to take advantage of 2D and 3D acceleration to improve basic display performance.

of the screen, and choose "View → Virtual Screen 1 → Use Host Screen 2" to move it to your second display, where it should scale up properly according to your previous "Scale Factor" settings (as an aside, you can reconfigure these settings from the same "View → Virtual Screen x" menu if they're incorrect).

» You should find your secondary display starts up automatically with the first, whether displayed on a secondary monitor or in a virtual window (Image B). If it doesn't, check for an option to enable it under the "View → Virtual Screen 2" menu.

## 3 NETWORK SETUP

Typically, when playing with your VM's network settings, you'll choose one of two "Attached to:" options: "NAT" for Internet access only, or "Bridged" to connect your VM to your network. Sometimes, however, neither may be desirable, but thankfully VirtualBox makes it possible to choose more flexible setups with a range of other options.

» Let's start with "Internal Networking," which enables you to join two or more VMs in their own private network while locking out outside devices. By default, this internal network

is called "intnet," but you can change it if you wish (just be sure to change the name across all your VMs). If you'd like to include your host PC (but no other network devices) in that network, choose "Host-only Adapter," which creates a loopback interface to join your VMs to the host, but no one else. For slightly more access to the outside world (for example, to make it possible for other network devices to access virtual servers indirectly), try "NAT Network."

» Both "NAT Network" and "Host-only Adapter" require additional global configuration before you can use them. Visit "File → Host Network Manager" to set up Ethernet adapters (complete with IPv4/IPv6 address masks for your VMs and PC) before enabling "Host-only Adapter," and "File → Preferences → Network" to set up NAT networks for your VMs to join.

» If you want your other network devices to connect to a virtual server on your newly created NAT network, start the VM you've assigned to it, and make a note of the IP address it's been allocated within the VM. From

## COMMAND-LINE USAGE

The VirtualBox GUI is a wonderful front end for VirtualBox, but it's not all-powerful. All available commands—and many more that aren't available from the GUI—can also be issued from the command prompt (Windows) or terminal (Linux), thanks to the VBoxManage command.

Before you can use VBoxManage in Windows, you need to register your VirtualBox directory as a path from which commands can be issued via the command prompt: Type "environment" into the Search box and click the "Edit the system environment variables" shortcut that pops up. Click "Environment Variables..." to bring up the Environment Variables dialog, then scroll down the "System variables" section, select "Path," then click "Edit." Click "New" and enter the path to VirtualBox, typically:

```
C:\Program Files\Oracle\VirtualBox
```

Click "OK" three times and you can now issue VBoxManage commands via the command prompt. The tool is covered in depth in Chapter 8 of the VirtualBox manual ([www.virtualbox.org/manual/ch08.html](http://www.virtualbox.org/manual/ch08.html)), but it follows this basic syntax: `VBoxManage <subcommand> "VM Name"`. For example, to start your virtual machine called "Windows 8.1" from the command line, issue the following:

```
$ VBoxManage startvm "Windows 8.1"
```

You can also create batch scripts in tools such as Notepad to automate multiple actions at once—see <https://medium.com/disruptive-labs> for details.

the VirtualBox "Preferences → Network" dialog, click the "Edit" button followed by "Port Forwarding." This enables you to route access to the virtual servers through your main PC—so, for example, if you'd set up a VM running BookSonic, you'd configure port forwarding to allow other network devices to access it via your main PC's IP address.

## 4 MORE FILE-SHARING OPTIONS

Navigate to a VM's "Settings" dialog and you'll see a "Shared Folders" option that enables you to create paths to specific folders on your host PC. Simply specify your host folder or drive, check "Auto-mount" for instant access, and enter a drive letter (such as T:) for Windows guests, or path `/media/shared` for Linux guests if you want to mount it to a specific point; otherwise leave it blank, and VirtualBox does the hard work for you. Your drive or folder appears like any other network share.

» Access should work straight away for Windows guests, but Linux guests may throw up permission errors. In that event, open a terminal window within the guest, and issue the following command, substituting `<user>` with your Linux username:

```
$ sudo usermod -aG vboxsf <user>
```

» Reboot your guest and you should now have access. The VirtualBox Guest Additions don't support older operating systems, so if you find you're unable to mount shared folders or connect to a network, another way to transfer files to the VM is to use VirtualBox's ad hoc VISO feature—basically a virtual ISO file that's mounted as a CD or DVD drive. With your VM running, go to "Devices → Optical Drives → Create ad hoc VISO." Use the top two panes to navigate to a specific folder, then drag files and folders into the bottom pane ([Image C](#)). Click "OK" and a virtual drive with your selected files appears in the VM. Sadly, you can't transfer files out of the VM from here. Also make sure you access the drive by right-clicking it and choosing "Open" in your Windows guest, because it may otherwise throw up an error about a missing application.

## 5 MAKE USE OF PASS-THROUGH DEVICES

Virtualization is, broadly speaking, the emulation of hardware through software, so even though it appears you're choosing hardware devices when picking an audio controller, network adapter, or graphics controller, they're actually software-based. This means a trade-off in performance, and restricts you to certain types of devices, but on the upside, they've been chosen to maximize compatibility across different operating systems (and VirtualBox typically chooses different options depending on the guest you're configuring).

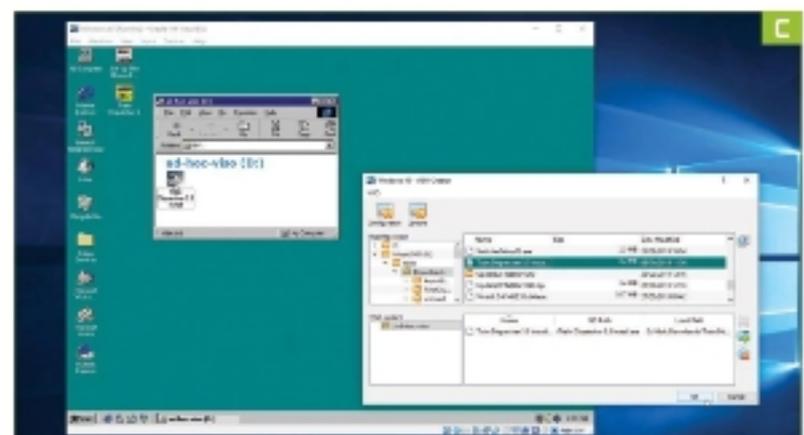
» VirtualBox is slowly developing more options to access underlying hardware. The USB option is one such example (although, even here, the USB controllers are software-based), as is the option to choose your physical CD or DVD drive under "Settings → Storage," by selecting "Host Drive" from the drop-down menu. The USB option is particularly handy—add wireless support to your guest via a USB or Bluetooth dongle, for example, or plug in a controller to play some vintage games ([Image D](#)). Note that you need to source—and install—drivers for these in the usual way within your guest OS.

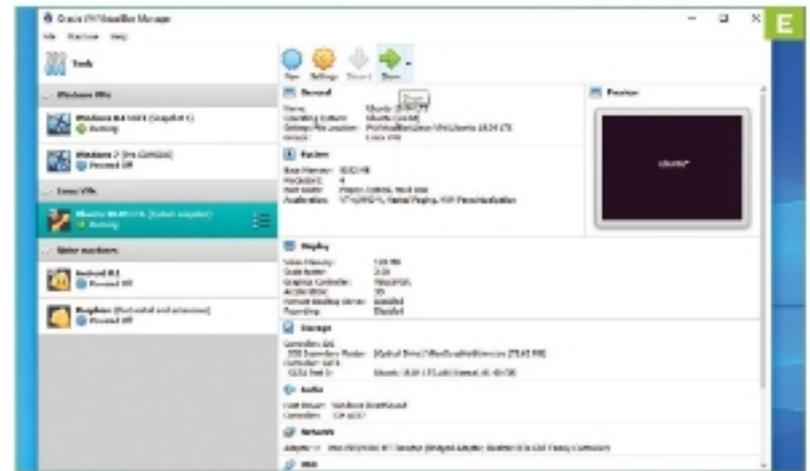
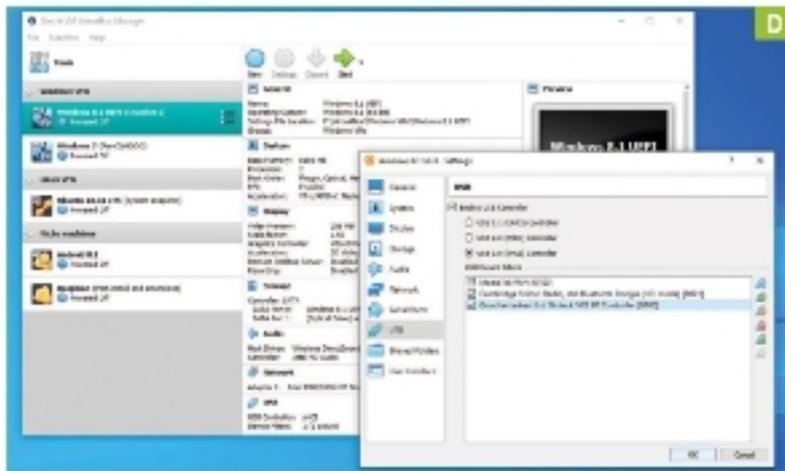
» Install the Oracle VM VirtualBox Extension Pack ([www.virtualbox.org/wiki/Downloads](http://www.virtualbox.org/wiki/Downloads)) and you gain a couple of additional hardware pass-through features to play with. The major one allows hosts to directly access any connected webcams by selecting them from the "Devices → Webcams" menu, while Linux hosts also gain theoretical access to an experimental PCI pass-through option that enables you to access regular PCI and some PCIe cards directly, too. Note: If you're looking to gain direct access to a graphics card on a Linux host with the thought of potentially running Windows games, it's a long shot at best. You might have more success experimenting with KVM ([www.linux-kvm.org](http://www.linux-kvm.org)) instead.

## 6 ACCESS DRIVES DIRECTLY

Another experimental feature makes it possible to attach entire physical partitions or drives to your VM, giving them direct access to their contents. This comes with a warning about data loss and requires mastery of the VBoxManage command-line tool (see box). If that doesn't put you off, and you accept responsibility for potentially wiping an entire drive, press Win-R, type "diskmgmt.msc" and hit Enter. Make a note of the disk number of the drive you wish to access, then issue the following via your host's command prompt or terminal:

```
$ VBoxManage internalcommands createrawvmdk
-filename "C:\Users\<user_name>\VirtualBox
VMs\<VM_folder_name>\<file_name>.vmdk" -rawdisk
\\.\PhysicalDrive#
```





» Substitute the path with where you wish to store the VMDK file, and give it a suitable name. Replace # at the end of `PhysicalDrive` with the disk number you recorded earlier—for example, `\\PhysicalDrive2`.

» This creates a small virtual VMDK file that points to your actual drive. You can then attach this as a data drive to your chosen virtual machine via its "Settings → Storage" page (click the floppy drive icon with the plus sign to add a new storage attachment). For more information—including how to access individual partitions only—see [www.virtualbox.org/manual/ch09.html#rawdisk](http://www.virtualbox.org/manual/ch09.html#rawdisk).

» A safer option might be to convert your physical drive to a VDI virtual image—it'll be larger, obviously, but your data will remain safe. Read on to discover how....

## 7 RESIZE YOUR VIRTUAL DRIVES

You know the drill: You created a VM with what you thought was adequate space at the time, only now it's no longer enough. Simple expansion of virtual drives is a straightforward task within VirtualBox itself: Just choose "File → Virtual Media Manager," select your drive under "Hard Disks," and then push the slider right, or enter your desired drive size (up to a maximum of 2TB).

» But what if you want to reduce a drive's size for any reason? VirtualBox prohibits this, so you need to head to <https://forums.virtualbox.org/viewtopic.php?t=22422> to download and run the portable CloneVDI tool. This works by cloning the virtual drive, so your original VDI is left untouched—simply point your VM to the new VDI file after it's been created in place of your existing one.

» To reduce the size of your VDI, simply select "Increase virtual drive size to" and enter a smaller amount (having checked

you've left enough free space at the end of the drive) into the box. CloneVDI also supports converting disk images to VDI from other virtualization formats, including Microsoft (VHD), VMWare (VMDK), and Parallels (HDD). You can also create a virtual disk image of a physical drive: Make sure you run CloneVDI as an administrator to gain sector-level access to the disk, then simply type `\\PhysicalDrive#` as outlined in the previous step in the "Source" box, and CloneVDI will detect it. Choose your destination, click "Proceed," and follow the prompts.

» More detailed instructions on using CloneVDI can be found in the "release notes.txt" file that accompanies the program.

## 8 HEADLESS AND DETACHABLE

If you're running a VM as a virtual server, you may not need to have a window into its operating system, in which case you can switch to a headless server by clicking the down arrow next to the "Start" button, and choosing "Headless Start." You'll see the virtual machine start up in the "Preview" window, but there will be no other indication it's running. If you subsequently need access to it, click the "Show" button to bring it into view [Image E]. Once done, click the close window button, then—when prompted—choose "Continue running in the background," and click "OK" to return to its previous headless state.

» VirtualBox also offers a "detachable" option—this is practically identical to headless, but starts up with a screen, so you can log in manually, for example, but you can close the window without shutting down. Again, clicking "Show" brings things back. ☺

## TRY NEW VIRTUAL MACHINES

Running most modern versions of Windows (7 onward) and Linux is simple in VirtualBox, but it's capable of emulating a wide range of older and niche OSes, too. Fancy dipping your toe into the Raspberry Pi's Raspbian desktop? Debian Buster with the RPI Desktop has just been released—download the ISO from [www.raspberrypi.org/downloads/](http://www.raspberrypi.org/downloads/)

[raspberry-pi-desktop/](http://raspberry-pi-desktop/) and follow the guide at [www.aokley.com/articles/2017-07-04-raspbian-x86-virtualbox.php](http://www.aokley.com/articles/2017-07-04-raspbian-x86-virtualbox.php).

How about emulating an Android tablet? Use the Android x86 ISO as a basis—visit [www.android-x86.org/documentation/virtualbox.html](http://www.android-x86.org/documentation/virtualbox.html) for a guide to follow.

Older Windows versions can be tricky, too, thanks in

part to a lack of support for the Guest Additions. A series of installation tutorials exists on the VirtualBox forums (look under "General → Howtos and Tutorials → Windows Guests"), but a simpler solution for Windows 95, 98, or Me is to visit <http://virtualldiskimages.weehiv.com/virtualbox.html> for a range of ready-made VDI images. And

speaking of Windows, heard about the ReactOS project? This attempt to render a lightweight form of Windows in an open-source environment is well developed, so why not give it a spin in VirtualBox? Get an ISO from <https://reactos.org/download>—try your older apps in this first before struggling with an older version of Windows.

# Monitor Tweaks for Competitive Gamers

## YOU'LL NEED THIS

### HIGH REFRESH RATE MONITOR

You need 120Hz or above.

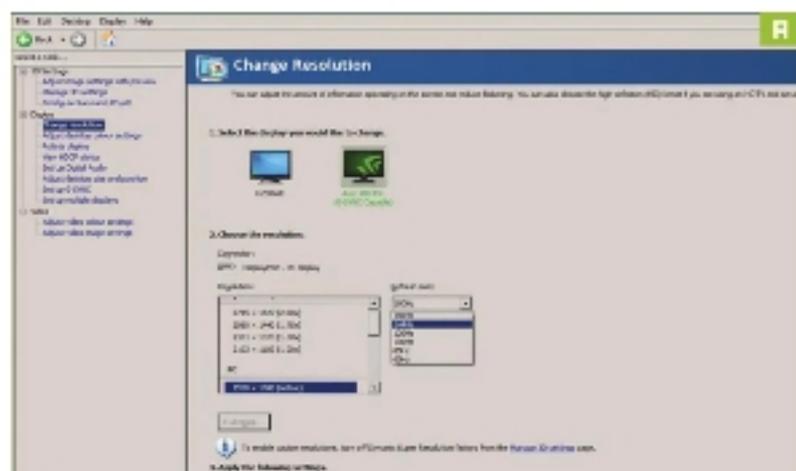
### NVIDIA CONTROL PANEL/ AMD RADEON SOFTWARE

Depending on your GPU, you need this to tweak and monitor performance.

## ASK A PRO PLAYER

about what's really important in a competitive gaming setup, and they'll all tell you the same thing: The truth hertz. As nice as it is to have RGB everything whirling away under your hands, and an RTX 2080 Ti churning out frames like tissue paper going through an industrial mulcher, it's refresh rate that often provides the biggest performance benefit of all. Improved reaction times, easier player movement prediction, and an overall sensation of smoothness with the mouse are all marginal gains that add up to something substantial at the top level. And because we all fancy ourselves as an undiscovered S1mple, we go out and buy the kind of componentry that the pros use, lost in the fantasy that every failure and subsequent teabagging in our online careers was simply the result of inferior hardware.

In the case of high refresh rate monitors, though, making the purchase alone isn't going to unlock the extra performance. Know what will? This handy how-to. There's an element of personal preference at play here, but the following steps will help you max out the advantage that high refresh rates afford you. —PHIL IWANTUK



## 1 SET YOUR REFRESH RATE IN WINDOWS

The new panel's out of its box, hastily assembled on to its stand, and an exciting new light glimmers forth. Excellent—but you're not going to feel the benefits just yet. Your first port of call is to right-click anywhere on the desktop, select "Nvidia Control Panel," then navigate to "Change resolution" under "Display," and check out the refresh rate option on the right [Image A]. If your previous monitor refreshed at 60Hz, Windows will probably still use this setting for the new screen, too, until you tell it otherwise. Do so by opening the drop-down menu and selecting the highest refresh rate available. You can also navigate to this option by finding your monitor in Windows Device Manager and hitting "Properties." You should see and feel a difference immediately when you move the cursor around on your monitor once the higher refresh rate is selected and applied.

## 2 OVERCLOCK YOUR MONITOR

Some monitors also include the ability to go beyond the purported refresh rate by "overclocking." No need to stand ready with the fire extinguisher, though—although there is a small chance of damage with any form of PC hardware overclocking, this is pretty tame stuff next to ramping up the core clock on your \$1,000 GPU. You need to enter your monitor's display menu (by pressing the buttons on your panel and waiting for the inevitably terrible menu to appear), then navigate to the overclock setting [Image B]. On the Acer Predator XB241H, for

example, there's an overclock menu within general system settings. If your monitor doesn't have a built-in method of overclocking, you can turn to software to up the hertz. Nvidia Control Panel, Radeon, and Intel Graphics Control Panel all feature custom resolution options that let you set a specific refresh rate manually.

» Note that the higher you go beyond your panel's factory settings, the higher the risk of frame-skipping or potentially damaging your hardware. And remember: There's no point running a refresh rate any higher than the frame rate you can consistently achieve in game.

## 3 DISABLE G-SYNC

This is a heartbreaker. You made the investment. You bought into G-Sync. And now you're going to turn it off. Why? Because although G-Sync has greatly reduced input lag versus regular old Vsync, it's still not as low as going without. And in games such as *CS:GO* and *Quake Champions*, where reaction time and precision are paramount, your aim is to minimize latency wherever it comes from. G-Sync makes things look nice and smooth, but that's not the same as being responsive. It's the difference between a luxury road car and a track-spec racing vehicle: One's going for comfort, the other for performance. After a few more steps, you might want to circle back and re-enable it, but we'll get to that.



## A NEW IMAGE

Why does refresh rate matter when we've been spending so many years (and dollars) chasing the FPS dragon? The answer is probably "marketing," because frame rate is only one half of the equation. Your GPU is the bottleneck in rendering new images to the screen, but your refresh rate is the bottleneck in how often those images are updated on the screen. So, if you have a 60Hz monitor and your game's running at 120fps, the screen's

still only being updated 60 times per second. And if you're gaming at 60fps on a 120Hz monitor, the same's true. Frame rates beyond 60fps do appear smoother at 60Hz, though, due to decreased input latency. It's just not showing a complete frame any quicker than 60 times per second. And that's where refresh rate comes in. At 120Hz and 120fps, your image is updated 120 times per second, with no wait between the GPU and monitor.



## 4 TEST YOUR PERFORMANCE IN GAME

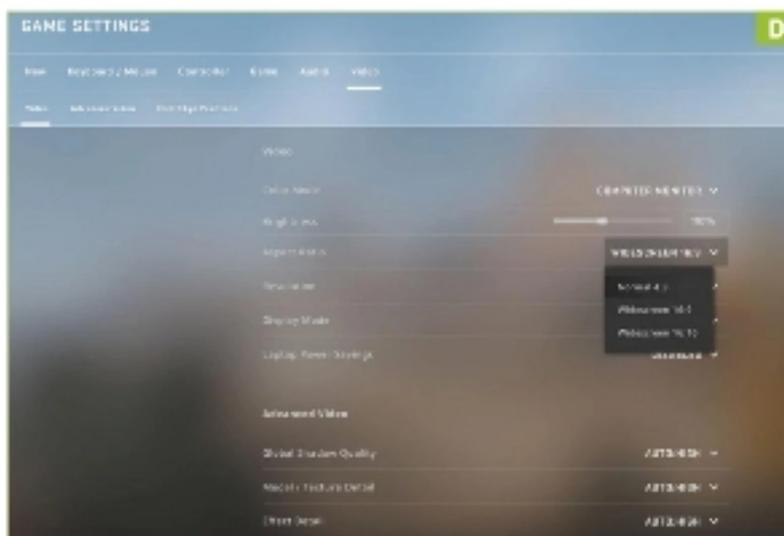
Now the important part. Load up your game, making sure Vsync is disabled in the display options to minimize latency, and just take its temperature. Play a few rounds with your new refresh rate [Image C]. How are you performing? Are you experiencing any stutters, noticeable input latency? Does your mouse sensitivity feel OK? It's possible that you might need to raise or lower the latter to suit the new smoother responses you're getting from the higher refresh rate. As for stutters and frame skips, these might creep in if you've overclocked your monitor a bit too far. Exit the game and revert back to the highest factory refresh rate available, then see whether you're still getting stutters. If you're not, you over-overclocked. If you are, it's likely a GPU/CPU issue.

## 5 LOWER SETTINGS SO THAT FPS MATCHES HZ

Here comes yet another reluctant moment—turning graphics settings down. Nobody likes to do it, but if you're going to make the most of, say, 144Hz, you need to be playing at least at 144fps, and a rock-solid 144fps at that. And it should go without saying at this point, but make sure Vsync is disabled. You can use Nvidia Control Panel's frame counter or in-game benchmarks to figure out your performance, then dial down the settings accordingly. Smoke effects and shadows can be the first up against the wall, along with any fancy lighting effects. If you're getting really desperate for performance, you can lower texture quality, too—just make sure the settings you're changing aren't making it more difficult to spot enemies in game.

## 6 EXPERIMENT WITH G-SYNC AND BRIGHTNESS

Now's the time to consider whether you'd rather have the smoothness and zero screen-tearing that G-Sync offers, or the responsiveness and minimal latency of going without. In



truth, this comes down to a combination of your chosen game and a splash of personal preference. None of the *CS:GO* or *PUBG* pros are sitting there with G-Sync enabled, but in a title like *Dota 2* or *Starcraft II*, where you're not aiming with pinpoint precision or moving the camera around with mouse freelook, the latency might not make much difference to you. If not: hooray! All that extra money you paid on that monitor wasn't for nothing! If so, well, at least you've got your K:D to look at.

Now's the time to make your experience even uglier by messing around with in-game brightness and raising it until enemies in darker areas are clearly visible, but not so much that you lose the definition between objects and surfaces. No, it doesn't look nice. But if you're going to go the whole hog on this, ignore the "barely visible" configuration image.

## 7 SWITCH TO 4:3 RATIO

A final rub of salt in the wound for the aesthetes: Let's stretch your aspect ratio out so everything looks fat and wrong [Image D]. It may seem like a weird move at first, but the majority of *CS:GO* players use a 4:3 aspect ratio on 16:9 screens, for two reasons. Firstly, running at 1280x900 on a modern machine gives you tons of frame rate to play with, and arguably even more importantly, the stretched aspect ratio makes your targets bigger on your screen. A sneaky move, but one that pays dividends when it matters. And with your game now thoroughly ruined in the eyes of the casual observer, you're ready to compete without the slightest excuse that your monitor's letting you down. ☹️

# Get to Know the New Luminar 4

## YOU'LL NEED THIS

### LUMINAR 4

Download a free trial from <http://skylum.com/luminar>

**LUMINAR 4 ARRIVED LAST MONTH**, and sees Skylum's image-editing app really come into its own with new "AI" cleverness and a general overhaul of its features. It still resembles Lightroom more than Photoshop, with its ability to display every image in a folder and let you take your pick of which to edit. It's non-destructive, so you never save over your original file, instead exporting a new one that represents your changes.

The AI features are particularly interesting, providing image analysis and improvement options that can be applied as a one-shot "Look" or as an "Edit" that gives you more control over the outcome. There's everything from all-over image enhancement to smart sky replacement and portrait-centered tools.

A free trial of Luminar can be obtained from <http://skylum.com/luminar> (drop down the Luminar menu at the top), where you'll also find a full beginner's guide. We produced this tutorial using a pre-release version of Luminar 4, and while the basic functionality will be the same, some interface features may look different on the final release. —IAN EVENDEN



## 1 MEET THE INTERFACE

For all its powerful tools, the Luminar interface [Image A] takes some getting used to. It's determined to do things its way, which means unlike any other app. It also doesn't believe in tooltips. Hover your mouse over something for as long as you like, and no snippet of text is coming to your aid. The usual menus are in the usual places, and "Undo" still lives on the "Edit" menu, so you can press unfamiliar buttons with no fear of something going badly wrong. Some useful tools are spread out across the top of the interface, such as the Eye that reverts the view back to your original image. Next to the Eye is a split-view button that provides a before and after representation of your image, with a draggable divider you can position to see the effect of your edits. "Crop" and "Export" have their own buttons, and the "+" at the far-left gives import options, of which there are two: import a whole folder of images, choosing from a filmstrip that appears on the left, or import individual images, aka Quick Edits.

## 2 START EDITING

Once you've got an image, there's a couple of ways to proceed. The quickest is to use one of the Looks, one-shot filters that live in a filmstrip of thumbnails at the bottom. One of the first ports of call, if you're looking for a quick tune-up of your image, is the AI Image Enhancer, in the "Essentials" section, which analyzes your image, and applies color and contrast edits to it automatically. It's not perfect, but it's better than many one-

shot filters. If the effect is too strong, every Look has a slider that can tone it down [Image B].

## 3 OTHER OPTIONS

If you're editing a landscape, there may be more appropriate options. Many types of photography have their own set of Looks, and there's an AI Landscape Enhancer that's more targeted to rolling hills and blue skies. We're also fond of Warm Sunset and Autumn Colors, which add a touch of the golden hour to an image.

## 4 LAYERS

Looks are one-shot: You click them, choose their strength, but there's nothing else you can do. If you want to build up a series of Looks, however, you can do it using Adjustment Layers. The "Layers" button is found all the way over to the right of the interface, and looks like a stack of squares. Using the "+" button here enables you to add an Adjustment Layer on top of your current image, which lets you use another Look. You can use any of the Edits, too, which are found under the pencil and ruler button under the "Layers" one. We quickly made an old-fashioned almost-sepia look by placing the Landscape B&W Look over the top of the Warm Sunset Look, and fading the B&W back with the Opacity slider using the Lighten Blend Mode [Image C].





## 5 EDITS

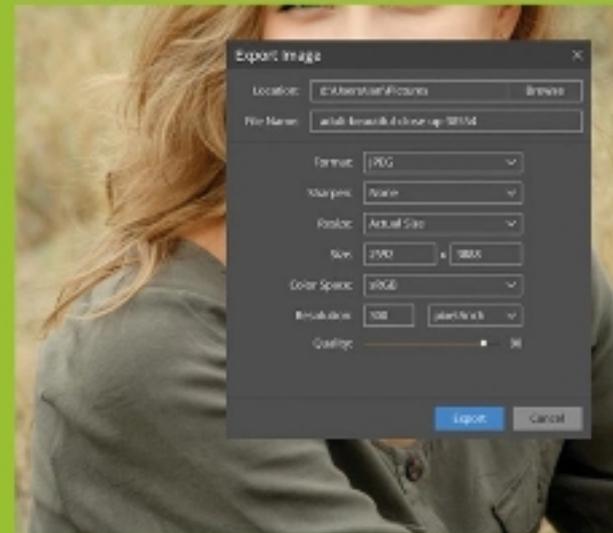
Opening up the Edits panel initially gives you some basic crop, rotate, erase (a bit like Photoshop's context-aware fill, but not quite as advanced, still clever and effective, though), and clone stamp tools. To get to the real meat of the Edits, you need to press another button—any of those at the far-right. Luminar's insistence on doing things its own way can be frustrating at first, as it can make it hard to find what you want, but with practice, it all becomes familiar. Edits are split into Essentials, Creative, Portrait, and Pro, with some old tools at the bottom in a set called Deprecated. For our landscape, we can use any of Essentials, but in Creative, there's a real star waiting.

## 6 SKY REPLACEMENT

Luminar 4 can recognize the sky in your image, mask off everything else, and either work on just that area or replace it completely. It starts in Essentials, with AI Sky Enhancer under "AI Enhance," but the AI Sky Enhancer in Creative is really good fun. You can choose from a built-in selection of cloudy, dramatic, or sunset skies, or use your own. The app removes the sky from your image and replaces it with the one of your choice, even going so far as to adjust the lighting in the rest of the image to match. It does a pretty good job, too [Image D], although it still gets caught up if there's something like grass or hair partly obscuring the horizon.

## 7 PORTRAITS

There are many tools in Luminar 4 aimed at enhancing portraits, including some face-detecting AI Edits to go with the one-shot Looks. There are the usual softening and smoothening options, including a Look named Female Portrait, which is disappointingly soft-focus compared to the rugged over-sharpening of Male Portrait. Elsewhere, you'll find some nice black and white options, too. In the Edits, there's an AI Skin Enhancer that offers the automatic removal of "defects" such as moles or scars, but otherwise just softens the skin. There's



## EXPORT OPTIONS

Luminar is a non-destructive editor, so your original file remains untouched by the process. You can import JPEGs or raw files produced by almost all digital cameras, and the ability to import folder structures means you can keep on top of all your images using the Library sidebar that appears in the same place as Edits. To export a file once you've finished editing, use either "File → Export" or the "Export" button awkwardly placed two-thirds of the way along the top bar. The "Export to File" dialog gives you options for what format you want, the level of compression, any additional sharpening you want, resizing, resolution, quality, and even which color space you'd like the file to use. It's a comprehensive list, and one that can be taken advantage of to produce a file suitable for almost any use.

Face Light, which cleverly just brightens the face in the image while leaving the background alone, and options for slimming the face, enlarging the eyes, "improving" (seems to mean darkening) the eyebrows and lips, as well as whitening the teeth. These head slightly too far into caricature territory sometimes, but the facial recognition tech that's behind their targeted edits seems sound [Image E]. Being non-destructive, every Edit and Layer can be toggled on or off using a switch on their respective panels, and Edits can be reset using an arrow button in the same place. ↻



# Push Live System Stats to the Cloud

## YOU'LL NEED THIS

### RASPBERRY PI

Any model is OK.

**TELEGRAM IS AN ENCRYPTED** private messaging service with an open API that is really easy to use. In this tutorial, we're building a Telegram bot—code that reacts to commands issued by the user from inside Telegram. Our bot will give us the latest CPU temperature, CPU usage, and current available RAM for a Raspberry Pi [Image A]. —LES FOUNDER



## SOFTWARE SETUP

Before any code can be written, we first need to sign up for a Telegram account [Image B]. There are many clients for Windows, Linux, and mobile devices; for this tutorial, it would be prudent to install the desktop client, but if you are completing the tutorial on a Pi, you can use the web version. With the account created and Telegram installed, we can now start to create a bot that reacts to messages. To do this, we need to use the BotFather, an automated script that handles the creation and maintenance of Telegram bots. Open a browser to <https://Telegram.me/botfather>, then tell the BotFather that we wish to create a new bot:

```
/newbot
```

Follow the automated instructions and after a short while your new bot is created. The BotFather provides a URL to start our bot, and an API key. Copy the API key somewhere safe—we'll need this later. Click the link to start a conversation with your new bot. You'll see a "Start" button at the bottom of the screen—click this to start the bot. That is all we need to do in Telegram for now, but we will come back to it later.

On our Raspberry Pi, we now need to install the Python library for this tutorial. Open a terminal and type:

```
$ sudo pip3 install pyTelegramBotAPI
```

Close the terminal when complete.

## 2 CODE THE PROJECT

Open your preferred Python editor (IDLE, Thonny, Mu, or similar) and create a new file. Save the file as "TelegramBot.py" and remember to save often.

We start the code for this project by first importing libraries that will provide the functionality in our application. The first is the "telegram" library for Python, then we import "time" to control the pace of the code. Finally, we import a library to enable Python to investigate system resources.

```
import telebot
```

```
import time
```

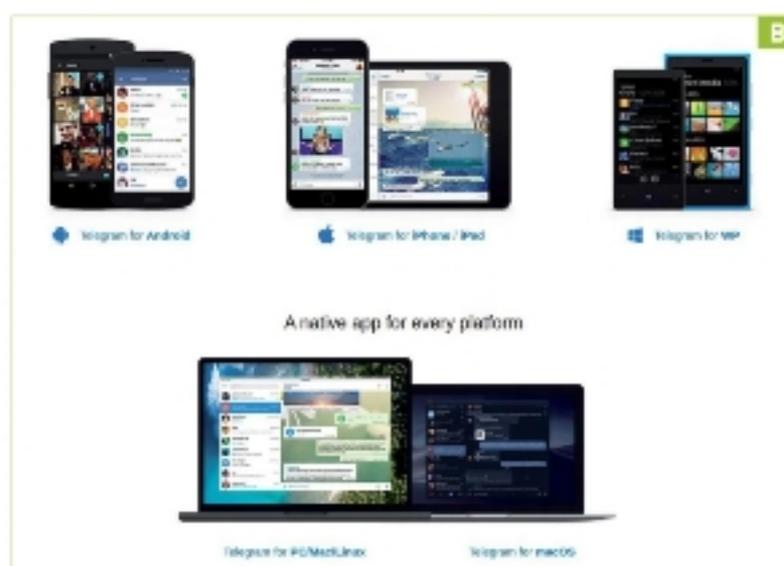
```
import psutil
```

The next step is to create an object that will provide a link for our code to the Telegram bot. This is handled via the API key, unique to our bot. The object `tb` creates the link:

```
tb = telebot.TeleBot("Your API KEY HERE")
```

With a connection to Telegram created, we now create a handler that reacts to a specific message from the user. The message is `/stats` and in order for it to be identified by the code, we need to create a test. The test is in a function that checks the contents of the message [`msg.text`] to ensure that it is not blank, and that it contains the `/stats` command:

```
@tb.message_handler(func=lambda msg: msg.text is not None and '/stats' in msg.text)
```





### 3 ACTIVATE!

Next is the creation of a function, `send_welcome`, that triggers the start of a flurry of activity. First a variable `temperature` is created and in there we store the output of the command to read all of the sensors on the Pi. But the output from this command is quite densely packed, so on the next line we update the contents of the variable so that it targets only the CPU temperature sensor. For the final step, we convert the returned variable into a list, so that the contents can be easily selected.

```
def send_welcome(message):
    temperature = psutil.sensors_temperatures()
    temperature = (temperature['cpu-thermal'])[0]
    temperature = list(temperature)
```

» As well as showing the CPU temp, we also want to show the CPU usage as a percentage, and the amount of free RAM available. It's just one line of code to get the CPU usage and store it in a variable called `cpu`, but for the memory, we first need to save all of the memory information to a variable called `memory`, then extract the current available memory, which is stored in a list with an index of 1. This makes it the second item in the list. Then we use some math to convert bytes to megabytes, and the `round()` command to produce an easily readable value.

```
cpu = psutil.cpu_percent(interval=1)
memory = psutil.virtual_memory()
memory = memory[1]
memory = round(memory / 1024 / 1024)
```

» Still inside the function, the next three lines handle replying to the user with the details of the server. First is the CPU temperature. Using the `tb.reply_to` function, we reply to the original message via a string. The string contains a sentence that tells the user what the data is, and inserted into that sentence is the contents of the list `temperature` that we created earlier. But we only need a specific piece of information, which is stored in index 1 of the list—this is the CPU temperature in Celsius. At the end of the sentence, we add the letter "C" to identify that we are using the Celsius temperature scale.

```
tb.reply_to(message, 'CPU Temperature is '+str(temperature[1])+'C')
```

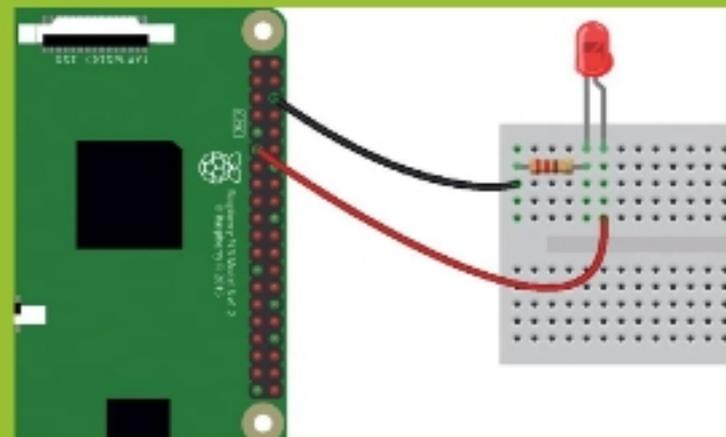
» The next two lines provide the data for the CPU usage and the amount of free memory. They work in a similar manner to the temperature line.

```
tb.reply_to(message, 'CPU Usage is '+str(cpu)+'%')
tb.reply_to(message, 'Free RAM is '+str(memory)+'MB')
```

### 4 FINAL CODE

This is the end of the function, so we can now move on to the final section of code, which is a loop to continually

## CONTROL THE GPIO



In the main tutorial, we created a function that would react to a command from the user, and post data to Telegram. We can also use Telegram to control the GPIO of a Raspberry Pi.

The code remains very similar to the main tutorial code, but we use the "GPIO Zero" library, and import the LED class, then we set GPIO 17 as the pin for our LED.

```
from gpiozero import LED
led = LED(17)
```

We then change the message handler to look for `/led-on`. If that command is issued, it runs the GPIO Zero code to turn on the LED.

```
@tb.message_handler(func=lambda msg: msg.text is not None and '/led-on' in msg.text)
def send_welcome(message):
    tb.reply_to(message, 'LED On')
    led.on()
```

The code to turn off the LED is very similar to turning it on:

```
@tb.message_handler(func=lambda msg: msg.text is not None and '/led-off' in msg.text)
def send_welcome(message):
    tb.reply_to(message, 'LED Off')
    led.off()
```

Run the code and you can now control an LED over Telegram! We've made a video demonstrating this—see <http://bit.ly/xf254telegram2>

check Telegram for updates | polling |, and to handle any error messages, and the user exiting the code. Using exception handling, we first try polling Telegram, but if the user requests to exit, an exception is raised and the code ends. Another exception is when an error is raised—this typically happens after an hour of use. If so, the code pauses for 15 seconds before the loop repeats.

```
while True:
    try:
        tb.polling()
    except KeyboardInterrupt:
        print("EXIT")
        break
    except Exception:
        time.sleep(15)
```

» The code is now complete and is ready to run. Start the code, then go to the Telegram application, and send the `/stats` message to your bot (Image C). This triggers the bot code to run and gives us live stats from our Raspberry Pi. It's as easy as that! 🏠

# Edit Video for Free

## YOU'LL NEED THIS

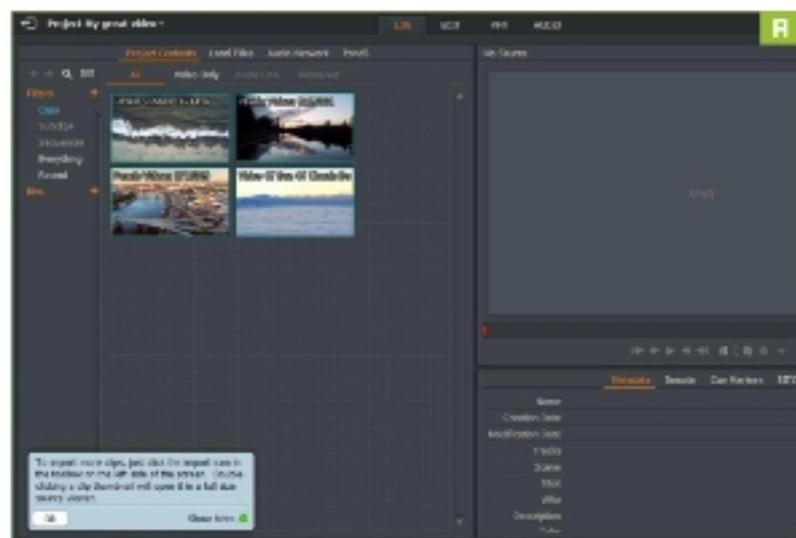
### LIGHTWORKS

Download the free software from [www.lwks.com](http://www.lwks.com).

**WOULD YOU LIKE** to use the same video-editing software as your favorite film-makers? Lightworks has been used on *Pulp Fiction*, *Heat*, *The Wolf of Wall Street*, and many more—and you can use it for your own videos, too.

The core app is free, but if you're a pro, you'll want to shell out for extra tools and the highest resolution exporting; for most of us, though, free is just great. The app contains everything you need to edit movies and export to common file formats and services such as YouTube.

Before you decide to make *Star Wars* in the shed, a word of warning: Video editing is very demanding on your PC, so it might be wise to keep things simple while you're learning the ropes. And don't let other programs run when you're editing video—your app needs as much power as your PC can give it. —CARRIE MARSHALL



## 1 CREATE A NEW PROJECT

When you launch Lightworks, you're taken to the "Local Projects" screen. Once you've worked on some projects, the various rectangles are replaced by their titles and thumbnails, but to begin with, everything is empty. The app doesn't have the familiar Windows menu bar—for example, the settings bit is accessed by the gears icon in the top-right corner. Click "Create New Project" and you'll see the "New Project Details" pop-up. You need to give your project a title, and the app also asks you to choose a frame rate—this is how many frames per second the video will be. If you don't have a specific one in mind, just stick with the "Auto" option. If you wish, you can also add quick notes in the text field to describe your project.

## 2 FIND YOUR FILES

If you're new to Lightworks, it provides on-screen tips, but you can turn them off. The Log view is where you locate and organize the source material for your project: your video clips and your audio tracks. When you're working on a new project, the "Project Contents" bit is blank. Click "Local Files" to browse to and import the video clips that you want to use. One of the key reasons for choosing Lightworks is that, unlike many video editors, it doesn't just work with one or two video file formats; it can import pretty much anything, which is a real boon if you're pulling in video from compact cameras or some smartphones. Simply browse to the folder containing your video, and select the clip(s) you plan to use.

## 3 IMPORT THE CLIPS

It's not immediately obvious, but the videos in your "Local Files" panel aren't added to your project until you import them. To do that, select them, then click the "Import" button in the

lower left-hand corner [Image A]. This can sometimes make your PC look as though it's crashed, especially on lower-powered PCs, but be patient—the process can take a few minutes, especially with long, hi-res video. You'll now see a three-panel view: clips in the top-left, preview at the top-right, and the project timeline at the bottom. Don't worry if you don't have any videos of your own to edit—you can do what we've done here, and download royalty-free videos from [www.pexels.com](http://www.pexels.com). These videos are available for you to use anywhere for free, including commercial use.

## 4 INSERT YOUR FIRST CLIP

To make your video, you need to put the component parts on the timeline. It's a little like Lego for video: You put one clip here, then another one alongside it, and another after that. Drag your first clip to the very beginning of the timeline. You'll see that Lightworks shows three things, not one, here: The bit on the V1 row is the video, while A1 and A2 are its left and right audio tracks [Image B].

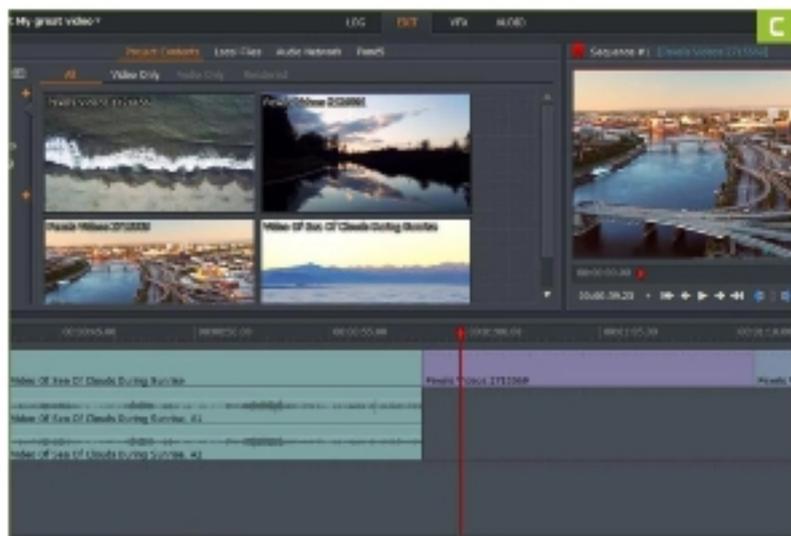
## 5 ADD ANOTHER

Most video projects are made of multiple clips joined together. Drag your second clip on to the timeline, just after the first one ends. When you get very close, it snaps to the end of the first clip, so there's no gap. Once again, we have three tracks here: one video and two audio. Lightworks color-codes each of your clips, so it's easy to see which one's which.

## 6 ADD A TRANSITION

Unless you're doing it for deliberate effect, it's generally a bad idea to move between very





different clips without some sort of transition, such as a fade or a wipe. To add a transition between clips, right-click the bit where the two clips meet each other, and you'll see the "Sequence" pop-up. Select "Transitions → Add," and choose the one you want. You can adjust its settings here, too.

## 7 ADD VIDEO WITHOUT AUDIO

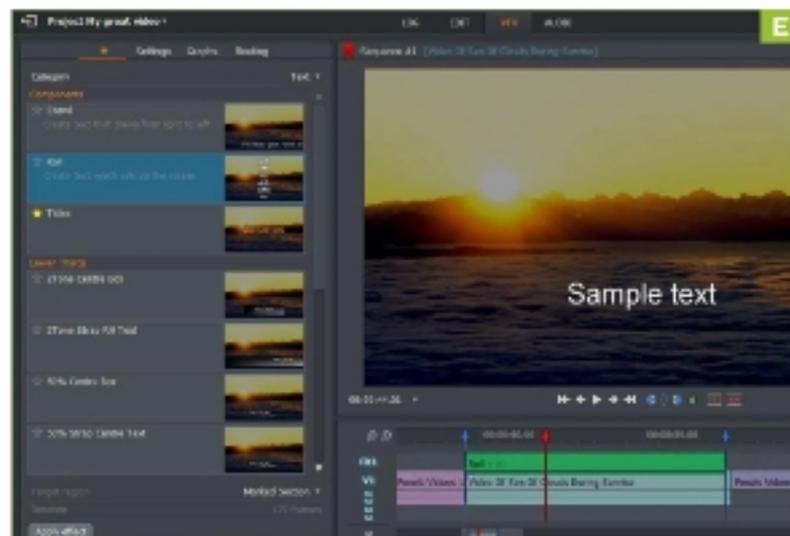
So far, all of our video clips have included their own audio, but sometimes your video will come without any audio at all. As you can see, two of the three clips here don't have audio tracks [Image C]. You can easily add audio by importing it and dragging it to the timeline, or you can click the microphone icon immediately below the preview (it's in the strip of small icons) to record a voice-over.

## 8 TRIM YOUR CLIPS

Trimming a clip gets rid of the beginning or end of it. To trim clips in Lightworks, right-click over the clip you want to trim, and you'll see the "Segment Commands" options. "Trim In" enables you to grab the edge and pull it in to get rid of bits you don't want without affecting other clips. Use "Slip" and "Slide" if you want to adjust where one clip ends and another begins.

## 9 ADD EFFECTS

If you click the "VFX" tab, you can do some interesting things to your video. The first time you click the tab, you're taken to the color correction screen. You can do two common tasks here: You can adjust the colors in a clip so it looks more natural or blends better with the rest of your project, or you can make more dramatic changes to its color. This is an extreme example of what happens when you change video settings [Image D]—we've whacked the brightness right down to give a more dramatic effect on our sunrise. Like any kind of filtering,



unless you're turning everything up to 11 for artistic reasons, it's usually best to do as little as possible when you color correct; too much can make your footage look distinctly odd.

## 10 ADD YOUR TITLES

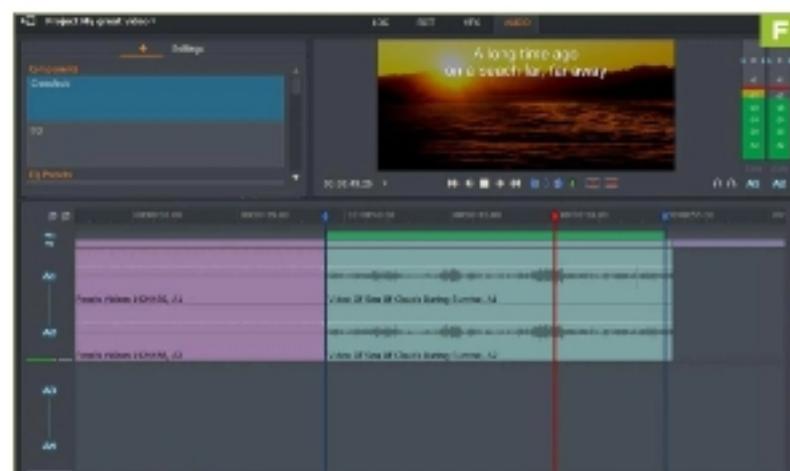
If you look closely above the color correction header, you'll see a little "+" sign. This is where you can add items such as scrolling titles or the video IDs you get on music video channels and YouTube clips. You can also get it by right-clicking a video clip in the V1 track and choosing "Edit → Effects → Add → Titles." We've gone for *Star Wars*-style text that rolls up the screen [Image E].

## 11 CHANGE THE TEXT

Your title settings appear in a new section underneath the "Color Correction" bit of the "VFX" window; scroll down to see all of it. Simply replace the sample text with whatever you want to see. You can also adjust the typeface size and opacity, its positioning, and its shadow or outline. Those two are particularly handy where the video moves from dark to light or vice versa.

## 12 ADJUST THE AUDIO

Click the "Audio" tab to add equalization or crossfades to your sound. There are some presets for reducing background noise, or you can roll your own, too. Adjust volume levels by dragging the blue arrows on the sound meters at the top-right [Image F]. To export your project, simply return to the "Edit" screen, right-click the timeline, and select "Sequence Commands → Export." 🎧



# Enhance Your Video Soundtracks

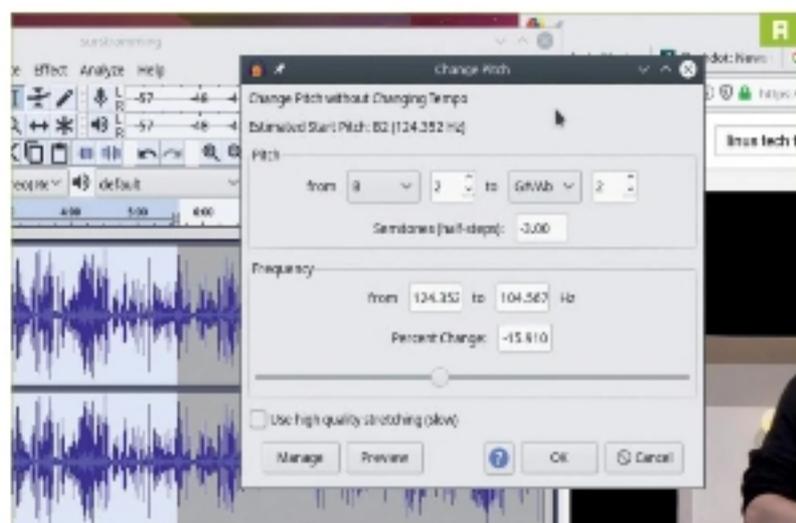
## YOU'LL NEED THIS

### AUDACITY

This cross-platform app is available for free from [www.audacityteam.org](http://www.audacityteam.org).

**WORRIED THAT YOUR VIDEOS SOUND AMATEURISH?** Is the sound dull, or the background noisy? Perhaps you're a YouTuber who wants to give your videos an edge, or a DIY film-maker seeking enhanced audio and sound effects. Before sinking time or money into big editing suites, check out what can be done on the easiest editor of them all: Audacity. It's available on Windows, Linux, and macOS, with a mature codebase and intuitive interface that has become something of a standard among audio enthusiasts.

If you want to tweak your video's sound, you need to extract it first, then put it back in the video when you're done, so we've included a short guide to this—see the box on the opposite page. We picked Kdenlive for our video editor, because it's by far the best known on Linux. For those just wanting some fun stuff to play with, we'll cover the sillier effects first, then move on to practical effects to help polish your audio. —JOHN KNIGHT



## 1 GETTING STARTED

In its default form, Audacity comes with a limited number of plugins under the "Effect" menu, which we'll examine first. Scroll down to the bottom of the list, and you will probably discover some "Plugin" submenus containing a list of extra effects. These were available on most installations we tried, but we can't guarantee their presence, so we'll cover these extra effects later.

➤ Most of the effects here follow the same flow of highlighting some audio, opening the "Effect" menu from the main menu, then choosing your effect. Usually, a new window opens with numerous settings and a "Preview" button, enabling you to adjust any necessary values before you change your audio. However, not all effects follow this workflow, and some require a little more explanation.

## 2 PITCH SHIFTING

Changing pitch enables you to tune audio up or down, so you can make someone sound like either a chipmunk or James Earl Jones. However, unlike the change of sound you get from speeding up or slowing down a record, pitch shifting does this without changing the tempo or length of the audio.

➤ Enable the effect by selecting your audio and choosing "Effect → Change Pitch." The new settings window is broken into two parts: "Pitch" and "Frequency" [Image A]. For the more tonally aware, "Pitch" is the place to start. Here you alter pitch

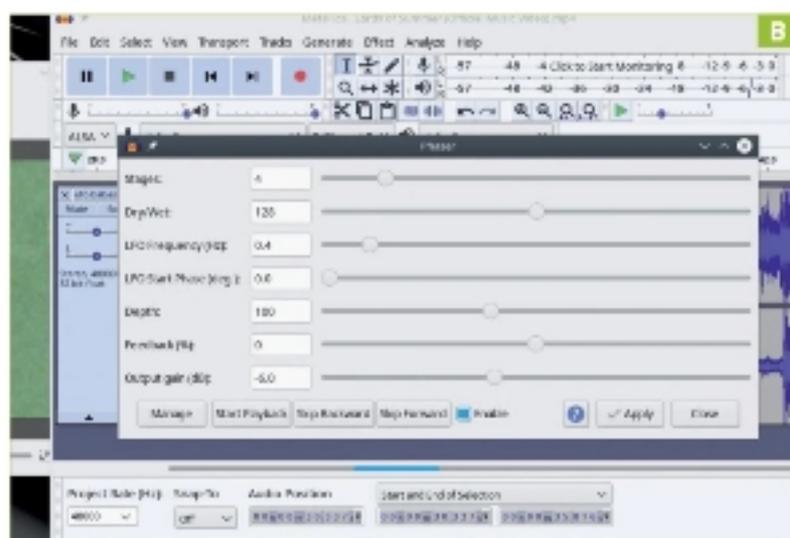
by starting from a musical note, and then choosing the note you want to shift into—whatever target you choose is displayed in the number of semitones below.

➤ If all of that is gobbledegook to you, just use the "Frequency" slider. Pull the slider right for something higher, and left for something deeper. A indicator tells you what frequency you'll shift from and to, as well as the percent of change taking place.

➤ Silly voices aside, pitch shifting has many practical uses. For instance, it can simulate an older voice, or how someone sounded in their youth. In the musical world, it is often used to compensate for changes in tempo and recording speeds, letting a recording be adjusted to fit another performance without changing any notes, or to adjust songs to another tuning.

## 3 WAH-WAH

This effect reproduces the sound of a wah-wah pedal. Famous the world over since the 1960s, this sound is used mostly by guitarists, where the signal phases in and out to create a funky, trippy sound. Perhaps the most well-known example is the famous guitar riff from "Voodoo Child (Slight Return)" by Jimi Hendrix. In a cinematic setting, a wah-wah effect could be used in cheap movies for alien invasions or a trippy dream sequence. To give it a try, select your portion of



sound and choose "Effect → Wahwah." Playing with the sliders in the "Wahwah" window, we found the "Resonance" control had the most immediately tangible effect, while the "LFT Frequency" slider changed the speed at which the effect phased in and out.

» Guitars and B-movie effects aside, wah-wah can add a new dimension to different instruments, such as during a bass solo (ask any Metallica fan), or to add a little something to vocal performances. For similar effects, try playing with any phaser or flange plugins [Image B].

## 4 REVERSE AUDIO

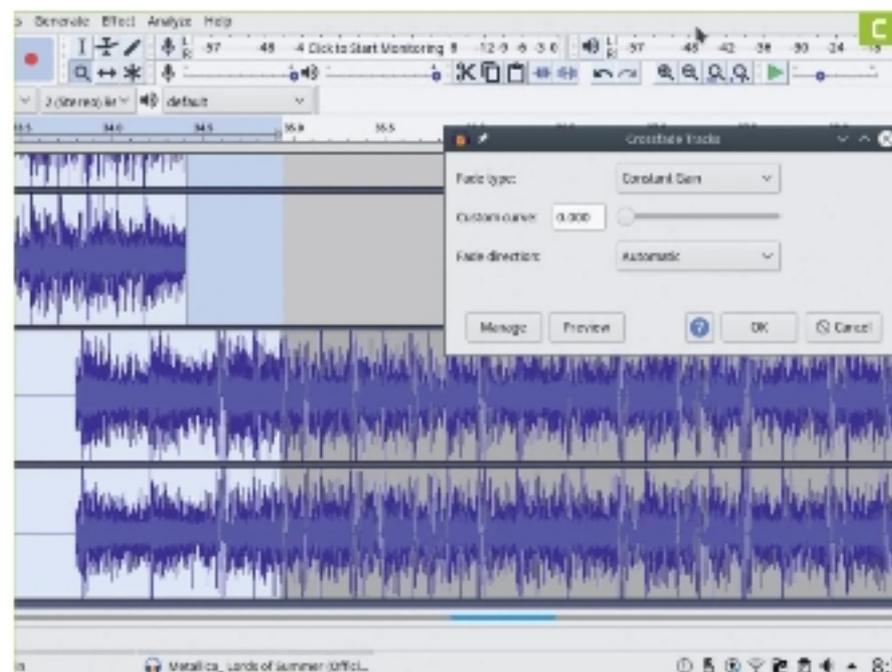
This effect is pretty obvious: Select some audio, click "Effect → Reverse," and that portion of audio is reversed. However, be careful where you place this in the waveform. If it's in the middle of a lump of frequencies, there could be a very rough jump in sound, along with a clicking or popping noise. Make sure that you either use this somewhere with some blank space either side, or see the later sections on fades and clicking and popping to clean up the sound.

## 5 REVERB

Adds that big resonant sound you get in stadiums, and is often built in to guitar amps and karaoke machines. If you've ever hit the reverb pedal on a piano, you'll know instantly the difference between a normal piano sound and the massive sound you get with reverb. Music aside, you can use reverb to simulate different environments. For instance, you can take spoken dialogue that sounded otherwise tiny and transform it into something that sounds as though you're in an enormous theater or aircraft hangar.

» To give it a try, select some audio, and choose "Effect → Reverb" from the main menu. In the new "Reverb" window, you can tweak all sorts of options like "Room Size," "Pre-delay," "Stereo Width," and so on. "Room Size" is fairly self explanatory: A value of 1 might sound like you're in a bathroom, while 100 may sound like a stadium. If you click the "Manage" button at the bottom-left, you'll find a number of factory presets, including vocal settings, differently sized rooms, and a cathedral.

» With creativity and experimentation, you can create some clever environment sounds. For instance, by turning down the "Room Size" and "Stereo Width," but cranking up the "Wet Gain," we were able to reproduce a very convincing sewer pipe environment. Used judiciously, reverb can transform your student DIY movies into something that feels more epic, or



## EXTRACTING AUDIO

After you've imported a video clip in Kdenlive ("Project → Add Clip"), extract the audio by choosing "Project → Extract Audio → Wav 48000Hz." This exports your current audio into lossless wave audio for editing with something like Audacity.

Once you've finished with Audacity, import your audio by clicking "Project → Add Clip" again. Any added clips are added to the "Project Bin" panel. These clips can be dragged on to the main editing field, and placed on any audio or video track you like—by default, a new track has multiple audio and video tracks ready to use.

To piece together your final product, first drag your video clip on to a video track, and mute the audio by clicking "Disable audio" in the control bar on the left of the track. Now add your newly mastered audio by dragging the audio clip from the Project Bin on to an audio track. Make sure both the audio and video tracks are properly aligned at 0:00, and your new movie should be ready to go. Try the preview window in the top-right to make sure everything has worked properly, and export your new video by clicking "Project → Render."

perhaps impart some stage presence to a performance otherwise devoid of charisma.

## 6 ADDING MORE EFFECTS

If you're keen to play with more effects, the Audacity website has a dedicated section for plugins, available on the Download page. Audacity has support for LADSPA, LV2, Nyquist, and VST plugins. Although LADSPA and LV2 plugins were originally designed for Linux, many of these effects have been ported to Mac and Windows. However, life is easier for Linux users, because your package manager is bound to come up with decent results, and any plugins should become available in your extended plugins menu.

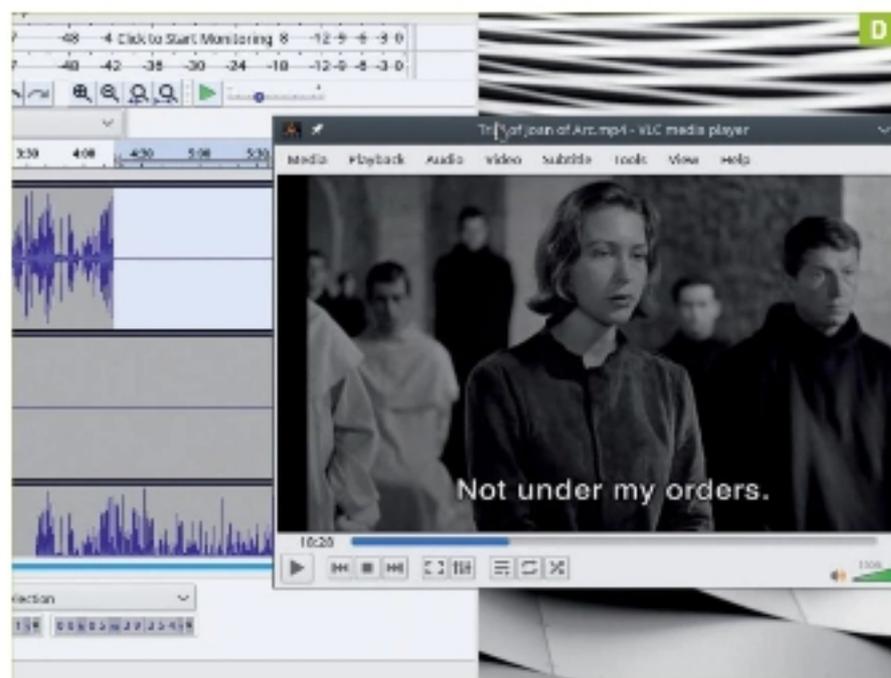
» If the plugins aren't in your menu, they may need to be enabled in the plugin manager, under "Effect → Add/Remove Plugins." Here, you can see everything available, whether it's "New," "Enabled," or "Disabled."

## 7 FADE CONTROL

Fades are easy to explain and use. They are what happen when a piece of audio starts or finishes with a gradual increase or decrease in volume. Rather than having an abrupt start or finish, it's a way of gently transitioning in and out of passages, making the experience more pleasant for the listener.

» Fades are very easy to do. To fade in, select from the beginning of your audio to the point you want full volume, and choose "Effect → Fade In" from the main menu. To fade out, select the audio from the point you want it to start decreasing until you want silence, and choose "Effect → Fade Out."

» Audacity's approach to fades is admittedly crude compared to something like Ardour, Qtractor, or Pro Tools—those apps have much more advanced functionality—but at least it's intuitive. If your fades are too drawn out or abrupt, it may be worth adding more silence to the beginning of the track and moving the fade



further up, or applying the fade more than once to adjust the curve at which the volume spikes.

» Not only are fades handy for transitions, but they are genuinely useful in cleaning up broken audio, as we will see later in this tutorial. Ultimately, fades are an essential part of any professional's toolkit, because they make your videos less jarring by avoiding abrupt halts in audio, as well as making smooth transitions and cool exits.

## 8 CROSSFADING

This one isn't in the main "Effects" menu, but it is among the extended plugins and is probably installed by default. This technique can be tricky to explain, but hopefully our picture [Image C] will clear things up. Crossfading takes two sound sources and slowly transitions between the two. In a multi-track editor like Audacity, this equates to having two tracks with a section of overlapping audio, then smoothly transitioning between the two.

» First you need to have your tracks sitting one above the other, so you can highlight the section of audio at the end of the first track, and the start of the audio in the second. Rather than having one track simply end and then the track below start, there needs to be some overlapping audio, so Audacity can fade between the two, so make sure each track has some "tail audio"—a couple of seconds should do. If you don't know how to move audio, select the Time Shift tool from the main toolbar or press F5.

» Now, with the select tool, highlight the tail audio of the top track, and while your mouse button remains pressed, drag the pointer down to the second track, highlight the first second or two of new audio, and release the mouse button. To apply the effect, choose "Effect → Plugins → Crossfade Tracks."

» Audacity now fades out the top track while simultaneously fading in the bottom track, with each track momentarily running over each other. Crossfading is very common in highlight reels and montages, and allows scene changes without bumps in audio or awkward silences—it is absolutely essential for DJs jumping between songs.

## 9 CLICKS AND POPS

If there are sudden loud clicking or popping sounds through your speakers, they're usually caused by very sudden changes in amplitude—often a result of badly pasted audio. There are two ways to fix this. First, you can try selecting the audio, and choosing "Effect → Click Removal." However, the

lower the threshold you set, the more likely it will have crackly side effects.

» For a more certain fix, find the bad part of audio, and zoom in closely on the wave. Highlight a portion of the end of the wave, and apply "Fade In" or "Fade Out," depending on which end of the audio you're fixing.

## 10 NOISE REDUCTION

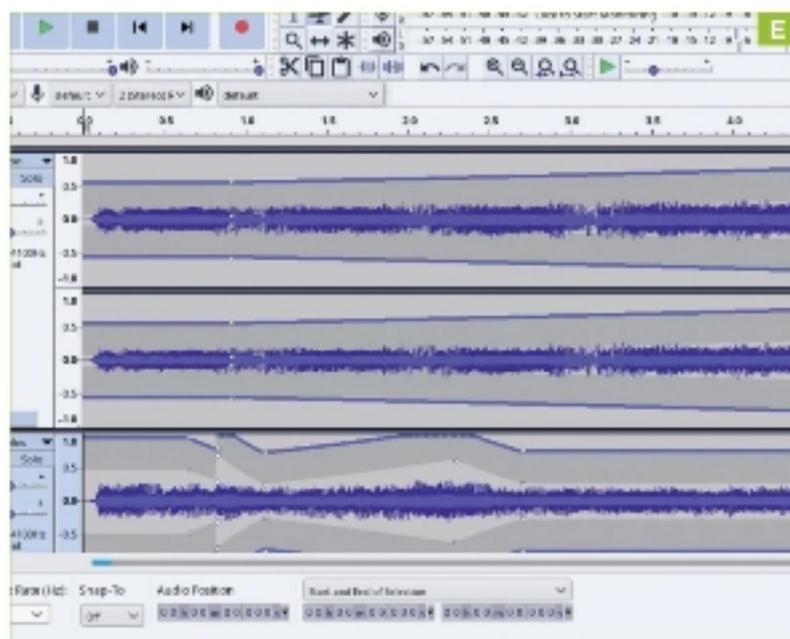
Audacity's noise reduction tool uses a different workflow to most plugins and can be initially confusing. First, Audacity needs a "noise profile" in order to analyze the track and know what to cut and what to leave alone. To get started, find and select a blank section of recorded audio that has background noise, such as hiss, ground noise, traffic, or whatever is bothering you. Now, from the main menu, choose "Effect → Noise Reduction." In the new window that appears, click the "Get Noise Profile" button. Audacity now drops back to the main window.

» Now choose all the audio you want to clean, which is likely to be the whole track. Open the Noise Reduction window again ("Effect → Noise Reduction"), and before touching anything, click "Preview" to see how the track sounds with the default settings.

» If the results aren't to your liking, start with the "Noise reduction (dB)" slider, which turns the effect up or down. It's tempting to just crank the setting way up, but the more noise reduction you apply, the more a weird "swimmy" sound will appear, along with very dull audio, much like a low-quality web rip. The trick is to strike a balance between removing enough noise for clarity, while still sounding natural.

» The "Sensitivity" slider broadens or narrows the number of frequencies let in. On some test audio near a busy road, 0 sensitivity let in all noise, but as we moved the slider right, first the traffic noise disappeared, then the bird noises went, and finally we were left with just dialogue—albeit very thin and unnatural-sounding.

» The last slider, "Frequency smoothing (bands)," is tricky to describe. Turned completely down, it had some very strange artifacts that sounded like sci-fi androids from the '50s were talking in the background. Turned all the way up, the dialogue sounded more natural, but with a strange phasing effect, almost like wah-wah. However, find the sweet spot, and it can really offset the "swimmy"



low-fi sound of too much noise reduction. If you're still not getting anywhere, at the bottom are two choices under "Noise"—"Reduce" and "Residue." "Reduce" is the default setting and is what you will have been using so far. However, switch over to "Residue," and you can hear what Audacity is cutting out. If any of the sound you're trying to keep is in the Residue audio, back off on the "Noise reduction" and "Sensitivity" settings.

» Ultimately, if you can strike the right balance between the three sliders, you should be able to improve the clarity of your audio, while adding a touch of professionalism.

## 11 BAD STEREO TO MONO

There is nothing more amateurish than a video playing sound through only one speaker, but fear not—mono is here to the rescue. You might think mono is a relic of yesteryears, but mono is perfect for certain applications: It's exactly centered, it comes through clearly on any sound setup, and it's half the file size.

» Mono is an excellent choice for simple spoken monologues (yes, that's you, vloggers), and any stereo music tracks can run alongside mono tracks in Audacity with no problem. For stereo tracks that have gone wrong, converting to mono centers the audio, and much of your audience won't even notice the lack of stereo effects, or even care if they do notice.

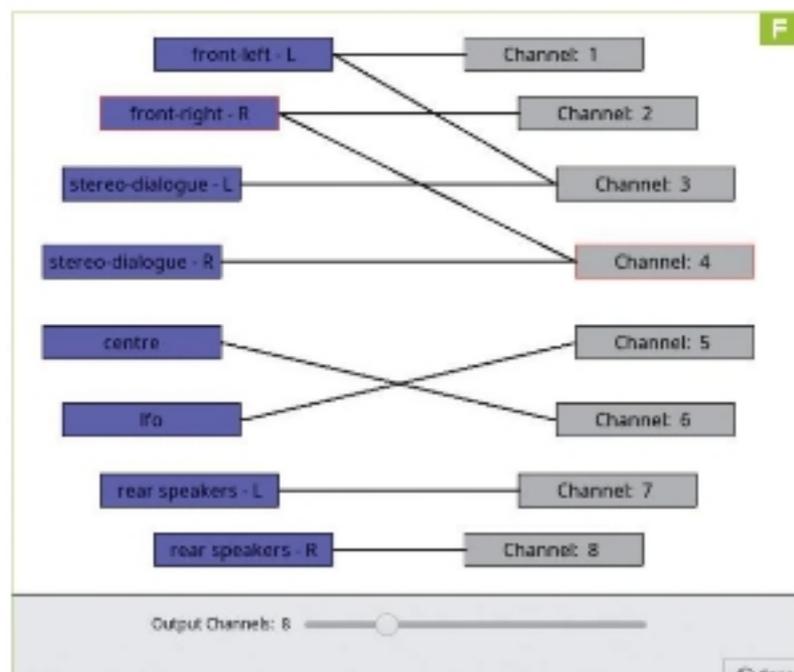
» There are two methods for mono conversion. For older versions of Audacity, click the track menu (left of the track), and choose "Split Stereo to Mono." You now have two separate mono tracks, one of which you can just delete. For newer releases of Audacity, from the main menu simply choose "Tracks → Mix → Mix Stereo Down to Mono."

» Going the other way, through clever use of isolation and pan controls, you can create new stereo soundtracks for old movies [Image D]—see the box on the right.

## 12 COMPRESSION AND NORMALIZATION

Do you have a soundtrack with inconsistent volume [Image E]? You can regain some consistency with compression and normalization. But what's the difference?

» While normalization raises sound levels across the track to something more consistent, compression takes a more aggressive approach to sound dynamics, making quiet parts louder and loud parts quieter. Normalization has less disdain from sound engineers as it still allows for some realistic dynamics, but compression's brute force approach is popular



## TAKE IT TO THE NEXT LEVEL

If you're frustrated with pure mic audio and want to make your soundtrack more flashy, you can use clever tricks to add positional sound. For instance, if you have multiple people taking turns to speak, or perhaps sound effects, such as tires squealing or gunshots, you can isolate these sounds by copying them into a new track, and muting their presence in the original source. Then, in the new track, you can simply use the pan slider to assign their position within the stereo landscape.

Techniques like these can even create stereo tracks out of old mono sound sources. This can get tricky when multiple sounds play at once, but clever use of EQ and techniques such as vocal isolation (or vocal removal) can help cut out other noises and focus on the sound you're trying to isolate.

The truly ambitious might try a surround sound conversion [Image F]. Audacity's interface isn't really built to handle surround sound—the GUI is designed around mono and stereo—but creating surround tracks is possible nonetheless. Although we won't be showing how to make surround mixes here (there are plenty of YouTube tutorials on this if you want to know more), we will give you one tip to get started.

Enable "Use Custom Mix" in the "Import/Export" tab of "Preferences," which opens a new mixing menu when you export your work. However, don't do this if you're only going to work in stereo, because it will greatly complicate your workflow. And, as always, remember to keep a backup version of your original audio!

with radio stations as it guarantees everything can be heard—very important if you have a weak radio signal.

» To normalize your audio, select your whole track by double-clicking the audio, and choose "Effect → Normalize." The option to normalize your maximum amplitude to -1.0dB is a good choice for dynamic music, but turn it up to 0dB if you don't care about dynamics.

» To compress your audio, select your whole track and choose "Effect → Compressor." We don't have space to fully cover compression, but the two most important options are "Make-up gain for - dB after compressing" and "Compress based on peaks." The first option cranks the volume right to the edge of the waveform, just before it distorts. "Compress based on peaks" can have weird side-effects depending on whether it's enabled. Use the "Preview" button to see whether you prefer it on or off.

» Be careful with compression—whispers sound like explosions, while explosions sound like whispers, and it makes everything sound like a beer ad.

» Once you've finished with your audio production, you need to export your work. Under "File → Export" there is a submenu where you can choose to export as MP3, WAV, or OGG—or just choose "Export Audio" if you need a different format. If you intend to share the audio online, use a compressed format, such as MP3 or OGG, but if you're about to use the audio in a video editor, use something lossless, such as WAV, as the final product will have compressed audio and video anyway. ⏻



# Subscribe from just \$26 a year



**Upgrade to a Digital & Print bundle for just \$10 more!**

**SEE THE WHOLE RANGE AND ORDER ONLINE AT:**

**[www.myfavouritemagazines.com/pcgsubus](http://www.myfavouritemagazines.com/pcgsubus)**

\*Terms and conditions: This is a US subscription offer only. Prices and savings quoted are compared to buying full-priced print or digital issues. Payment is non-refundable after the 14 day cancellation period unless exceptional circumstances apply. Your statutory rights are not affected. Prices correct at point of print and subject to change. For full terms and conditions please visit: [bit.ly/magandc](http://bit.ly/magandc)

# REVIEWS

TESTED. REVIEWED. VERDICTIZED.

## INSIDE

- 76 Falcon Northwest  
Talon 20th Anniversary  
Gaming PC
- 78 Asus ROG Mothership
- 80 Nvidia GeForce GTX  
1660 Super
- 83 AMD Ryzen 5 3400G
- 84 MSI MEG X570 Ace
- 86 TerraMaster F5-422
- 87 Noctua NH-D15 Chromax  
Black
- 88 Cooler Master MM830
- 89 Razer Seiren Emote
- 90 Borderlands 3
- 91 Tom Clancy's Ghost  
Recon: Breakpoint
- 92 Lab Notes

AMD  
RYZEN 5  
3400G  
PAGE 83



# Falcon Northwest Talon 20th Anniversary Gaming PC

This falcon is as beautiful as it is deadly

WHEN WE ASKED industry veteran Falcon Northwest to provide us with a desktop that could crush benchmarks and had a long-haired dude on the side, it delivered—no questions asked. Joking aside, the latest iteration of the company's popular Talon gaming PC is testament to FNW's excellence and experience in building awesome PCs. Falcon's towers are built for gamers and power users, and while you might want a slightly different configuration for video editing, our 20th Anniversary Talon was made to crush any gaming benchmark into a fine powder.

Falcon Northwest's Talon has base model configurations ranging from \$3,000 to \$4,000, and that's without checking any optional boxes. Our build came with all the trimmings for the tidy sum of \$5,316, and included a pair of RTX 2080 Supers, powered by a Ryzen 9 3900X CPU, with 32GB of G.Skill Trident Z 3,200MHz RAM. To top it all off, an AIO cooler emblazoned with the FNW logo kept the CPU temperature hovering around a steady 60 C, even under load.

Custom builds typically boast excellent cable management, but Falcon went beyond the call here—it was difficult to tell where all the cables were stashed, giving the interior of the case an eerily vacant look. While FNW is known for excellent build quality, its aesthetics have always been muted. Newer configurations, however, including the 20th Anniversary

version of the Talon, are available with either custom UV printing or a handful of pre-selected "Exotix" paint jobs. Our configuration sported a custom paint job, with a slick UV printed side panel featuring none other than Geralt of Rivia, whose muscular frame was merely an aperitif to the power contained within.

The side panels swing open and latch with a magnetic closure, so you don't have to unscrew anything to access the components. The front panel can also be quickly removed to expose the front dust cover. The front I/O features dual USB-A ports and a single USB-C port on the top of the case, along with a 3.5mm audio jack. The rear I/O is packed out with 11 USB-A ports and one more USB-C, courtesy of the Asus Crosshair VIII motherboard, which also provides integrated Wi-Fi.

The interior of the case has both GPUs in a traditional lateral mount, with an NVLink SLI bridge to concert their efforts. The Ryzen 3900X is cooled via an AIO 280mm radiator connected to a pair of front-mounted 140mm RGB fans. A single fan in the rear helps maintain airflow to the rest of the case. The case itself is a sturdy aluminum mid-tower that measures 17.25 x 8.75 x 15.5 inches, and weighs around 40lb. Not prohibitively large or heavy, but definitely not flimsy; the Talon feels like it was built to last.

And last. The Talon performed brilliantly against our top-tier zero-point,

matching or improving on the benchmark scores in all areas, although making the most of both cards in these post-SLI days is difficult. If you need to do 4K gaming at 60fps with all your sliders turned up, the Talon can provide that kind of power.

As with nearly all pre-built custom PCs, you're paying for a service, which can be a small price for not having to deal with things such as cable management or painting a buff man on the side of your computer. But it's important to know how much of a premium you're paying for the guts of your PC. All the Talon's parts, tallied up, cost around \$3,400, which means the custom case, three-year warranty, and subsequent assembly add around \$1,900. The total still beats similar custom builds from Origin.

FNW debuted the original Talon around 20 years ago. This model had the same excellent components and build quality, but was almost deceptively ordinary in its appearance. While you are paying a premium, the bottom line is that if you've got the cash, the Talon can provide you with uncompromising performance in a unique package. —ALICE NEWCOMBE-BEILL

VERDICT  
9

Falcon Northwest Talon 20th Anniversary Gaming PC

SOARING Awesome

performance; slick aesthetics; clean build.

BORING Expensive.

\$5,316, [www.falcon-nw.com](http://www.falcon-nw.com)

## BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	2,178	3,108 (143%)
CrystalDisk QD32 Sequential Read (MB/s)	3,136	5,587 (178%)
CrystalDisk QD32 Sequential Write (MB/s)	2,126	5,492 (258%)
3DMark: Fire Strike Ultra (Index)	6,988	12,551 (80%)
Total War: Warhammer II (fps)	42	75 (79%)
Tom Clancy's The Division 2 (fps)	38	42 (10%)
Metro Exodus (fps)	37	43 (16%)

Our zero-point consists of an Intel Core i9-7900X, 32GB G.Skill RipJaws V series DDR4 @ 3,000, an MSI GeForce GTX 1080 Ti Gaming X, and a 512GB Samsung 960 Pro PCIe SSD. All tests performed at 4K at the highest graphical profile.

## SPECIFICATIONS

Processor	AMD Ryzen 9 3900X
Graphics	2x Nvidia RTX 2080 Super
RAM	2x 16GB G.Skill Trident Z RGB 3,200MHz DDR4
Motherboard	Asus Crosshair VIII Hero
Storage	2x Samsung 970 EVO Plus 1TB SSD (RAID)
Cooling Solution	ASTK-680LS liquid cooler
PSU	EVGA 1000 G3 Supernova
Case	Talon with custom etching
Warranty	Three years parts/labor



The Talon can really get its claws stuck into anything you throw at it.

# Asus ROG Mothership

## Not a little green man in sight

RATHER APPROPRIATELY for a computer named after the lead vessel in an alien armada, the ROG Mothership looks like something from a science-fiction movie. More specifically, the bulky design and weight evoke an air of '90s comedy sci-fi, like something a grumpy Agent Kay might struggle to use while Will Smith hovers in the middle distance. Unlike the technology of the time, however, this Mothership is a computational powerhouse, packing everything from 64GB of memory to a ninth-gen Core i9 processor, and more.

Like a Microsoft Surface Pro that's been on a particularly glorious Burger King diet, this hulking laptop's keyboard folds down and detaches from the main body, but don't be fooled—this is no convertible machine. The keyboard doesn't fold around to the back of the unit, and while it does have a Surface-esque kickstand, the Mothership's glorious 4K display is not touch-controlled. This is a desktop replacement, for someone who needs to set up a powerful machine wherever they go, not a tablet for a train-riding digital artist. If you tried to set this beast up on your tray table for a bit of Skyrim on your transatlantic flight, you'd be swiftly ejected by some irritated fellow passengers.

In fact, it's a little unclear exactly who the target audience is for this product. Don't get us wrong, it's awesome, and we want five of them, but at \$10,000, it's difficult for even the flushest shopper to rationalize a purchase. Like a murderous alien that comes to appreciate the natural

beauty of humanity, the Mothership is having a mild identity crisis. The Republic of Gamers tag implies that this is strictly a gaming-focused system, as does the black-and-gold design, and flashy RGB lighting, but heaps of RAM, a 4K screen, and some (OK, somewhat dubious) portability sit it more within the purview of creative users. It's perfectly situated for connection to an external 4K monitor or two, making it effectively a great cornerstone for a home office setup.

Yes, it's great for gaming. With a GeForce RTX 2080 GPU, it's no slouch in the graphical department, although it does struggle at hitting the 4K potential that the display offers. Set it at 1080p or 1440p, though, and you'll be able to enjoy top-notch frame rates in just about any game. The keys are responsive, with a reasonable amount of travel, although a wireless gaming mouse is perhaps a necessary purchase, as the keyboard doesn't have a USB pass-through.

However, this system also works pretty well as a content creation rig. 64GB of memory is a respectable quantity for editing and rendering, and the Core i9-9980HK CPU is practically as good as it gets with laptop processors, comfortably able to handle CPU-bound tasks. Any video editor worth their salt would likely be pleased with the Mothership's specs, while a hardcore PC gamer might look at the price tag and decide that they're better off sticking with a custom-built desktop. There's plenty of overclocking potential here, at least; in fact, included in the box are two AC adapters, both

of which need to be plugged in for maximum overclocking.

Perhaps our most pertinent criticism falls on the design front. When attached to the main body, the keyboard's hinge is a little too flexible for our liking, meaning that the entire assembly shifts a little when typing. It doesn't feel like it's in any danger of breaking, but it was a minor annoyance that drove us toward using the detached configuration, which deactivates any RGB effects you've set up on the individually-addressable keys.

The only other beef we've got with this laptop—other than the fact that calling it a "laptop" is faintly ridiculous—is the kickstand. It's razor-thin, instantly reminiscent of the Surface kickstand. The problem is that the Surface is an ultralight quasi-tablet, while the Mothership is anything but. Here, the kickstand is tricky to set up, and feels like it could be easily damaged. These are minor annoyances, however; the Mothership is still an awesome powerhouse of a product, and the high price can just as easily be considered a premium for Asus's work in cramming all that tech into its chassis. —CHRISTIAN GUYTON

### VERDICT



### Asus ROG Mothership

SPACEFLIGHT Great for 1440p gaming; detachable keyboard works well.

CRASH LANDING Some flimsy design choices; extremely expensive.

\$10,000, [www.asus.com](http://www.asus.com)

### BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	1,030	1,821 (77%)
CrystalDisk QD32 Sequential Read (MB/s)	3,374	8,714 (158%)
CrystalDisk QD32 Sequential Write (MB/s)	2,530	4,636 (82%)
3DMark: Fire Strike (Index)	13,610	22,997 (69%)
Rise of the Tomb Raider (fps)	92	141 (53%)
Total War: Warhammer II (fps)	62	90 (45%)
Tom Clancy's Ghost Recon: Wildlands (fps)	49	74 (51%)

Best scores are in bold. Our gaming laptop zero-point is the Acer Predator Triton 500, with an Intel Core i7-9750H, 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 i9-9980HK
Graphics	Nvidia GeForce RTX 2080
RAM	64GB DDR4-2666 SDRAM
Screen	17.3-inch 4K UHD @ 60Hz
Storage	3x 512GB PCIe NVMe M.2 SSD
Ports	3x USB 3.1 Gen 2 Type-A, 1x USB 3.1 Gen 1 Type-A, 2x USB 3.1 Gen 2 Type-C, HDMI, RJ-45, SD card reader, audio
Connectivity	Gigabit Ethernet, 802.11ac, Bluetooth 5.0
Weight	10.6lb
Size	16.1 x 12.6 x 1.2 inches



The Mothership is cooled by two powerful fans and eight heat pipes.

# Nvidia GeForce GTX 1660 Super

## A good card to fill a narrow gap

**EARLIER THIS YEAR**, Nvidia launched its new GTX 16-series with the GTX 1660 Ti, GTX 1660, and GTX 1650 seeking to join the ranks of the best graphics cards. The lack of ray-tracing hardware is interesting after last year's RTX 20-series, but it makes sense for the mid-range and lower markets. When Nvidia refreshed the RTX lineup with the Super models, we didn't expect to see that extend into the GTX range, yet here we are. The GTX 1660 Super now exists, leveraging the same Turing architecture and TU116 GPU.

We're not saying the card is bad by any means, but it ends up overlapping the GTX 1660 Ti in many respects. It still has the GPU core counts and clock speeds of the GTX 1660, but it gets GDDR6 memory—which was one of the main differentiators between the 1660 and 1660 Ti. Even more curious, the GDDR6 memory is clocked at 14Gb/s, compared to the 12Gb/s speed on the 1660 Ti. Basically, it's weird to have a lower-tier card that's both better and worse than an existing higher-tier card. At least the price is lower.

The main specs are no surprise. It's the same TU116 GPU, still with 22 SMs enabled, now with 14Gb/s GDDR6. It's important to note that the 1660 Super

is intended to coexist with the other GTX 16-series GPUs, rather than replacing anything. That means Nvidia now has five GTX 16-series parts instead of three. The other interesting aspect about the GTX 16-series is the lack of hardware ray-tracing features. There's driver support for DirectX Raytracing, but generally speaking, you won't get much above 30fps even at 1080p. With the RTX 2060 now starting at around \$320, the GTX 1660 Ti is in a tough spot. With the 1660 Super's price of \$229, that's a big enough jump that the Super feels like a viable option.

Hopefully, this comes as a surprise to no one, but the GTX 1660 Super ends up performing just a bit slower than the GTX 1660 Ti in pretty much every test we ran. Overall, the 1660 Super is 3–5 percent slower than the 1660 Ti, and it's about 15–20 percent faster than the vanilla GTX 1660. The faster GDDR6 memory didn't provide the newcomer with any victories, but considering it has a 9 percent deficit in GPU cores, that's to be expected.

Practically speaking, there's almost no difference in real-world performance when looking at the 1660 Ti and 1660 Super. Short of running benchmarks, no one is going to be able to discern

a 5 percent difference, and tweaking one or two settings down a notch would make up for that deficit. Sure, a heavily overclocked 1660 Ti going up against a stock 1660 Super might show a larger gap, but a heavily overclocked 1660 Super should also come out slightly ahead of a 1660 Ti running reference clocks.

The more pertinent question is why is Nvidia releasing the 1660 Super? Does it fear AMD's RX 5500, or perhaps the lower prices on the RX 590? We don't know. Whatever the reason, the fact remains that the 1660 Super is better value overall. It's officially \$50 cheaper than a 1660 Ti—a difference of 18 percent—while performance is within 5 percent.

The thing is, while a \$50 difference in price isn't that big, particularly in the high-end and extreme markets, for mainstream and budget builds, where every cent counts, the GTX 1660 Super is an easy recommendation. It's nearly as fast as its big brother, it costs less, and it outperforms anything AMD currently offers in the same price range. And it does all that while using less power. If you can't go much over \$200 for a graphics card upgrade and you're still running a GPU that's several generations old, say hello to the GTX 1660 Super. —JARRED WALTON

### VERDICT

# 9

### Nvidia GeForce GTX 1660 Super

**■ SUPERCHARGED** Good 1080p performance; better value than 1660 Ti; proven Turing architecture.

**■ UNDERPOWERED** Missing RTX features; crowded price point.

\$229, <http://nvidia.com>

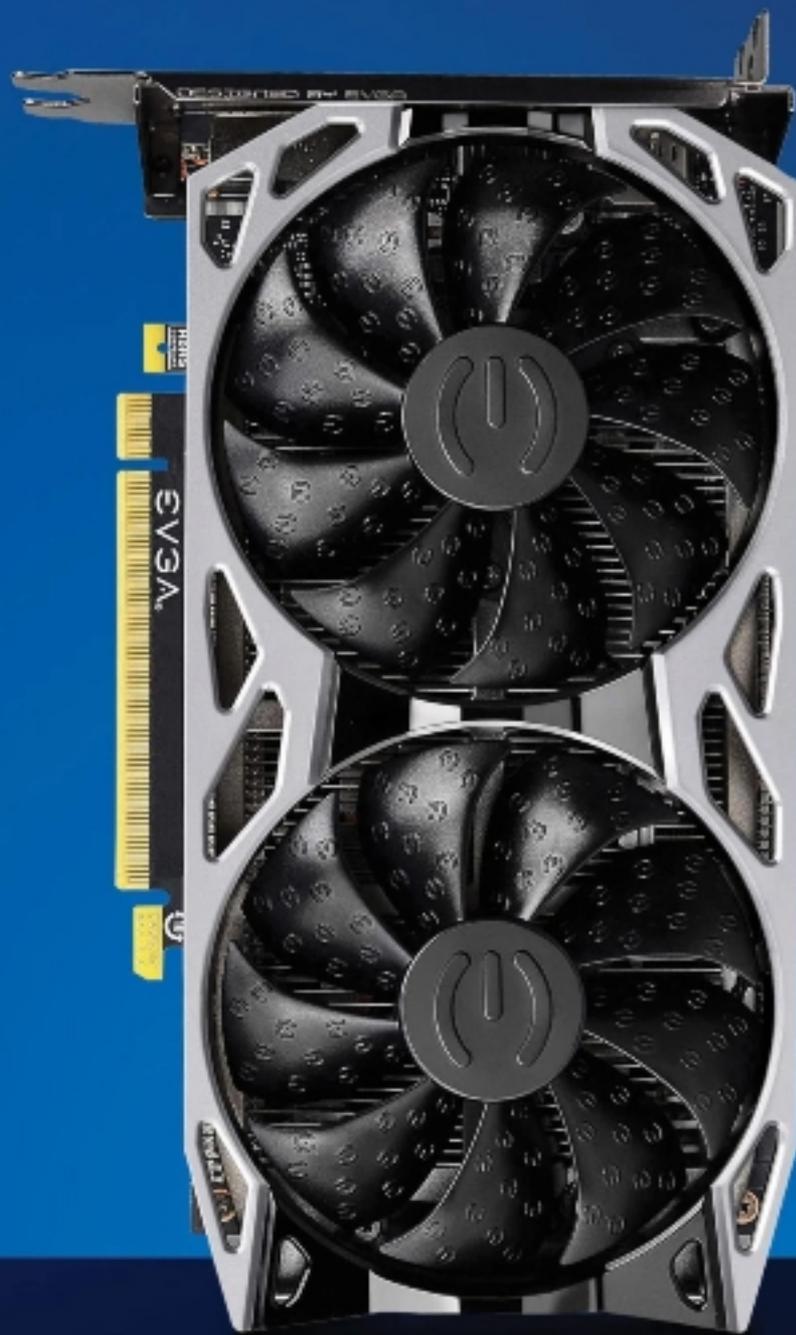
### BENCHMARKS

	Nvidia GeForce GTX 1660 Super	Nvidia GeForce GTX 1660
Tom Clancy's The Division 2 (Avg fps)	<b>70/45</b>	59/38
Assassin's Creed Odyssey (Avg fps)	<b>53/42</b>	46/37
Far Cry 5 (Avg fps)	<b>96/68</b>	81/57
Shadow of the Tomb Raider (Avg fps)	<b>85/56</b>	74/51
The Outer Worlds (Avg fps)	<b>81/54</b>	72/47
Metro Exodus (Avg fps)	<b>53/40</b>	47/35
Total War: Warhammer II (Avg fps)	<b>71/54</b>	63/46

Best scores are in bold. Our test bed consists of an Intel Core i7-8700K, 16GB of G.Skill DDR4-3200, a Gigabyte Z390 Aorus Gaming 7, and a 1TB Samsung Evo M.2 SSD. All games are tested at their highest graphical profile with AA turned on, at 1080p and 1440p respectively.

### SPECIFICATIONS

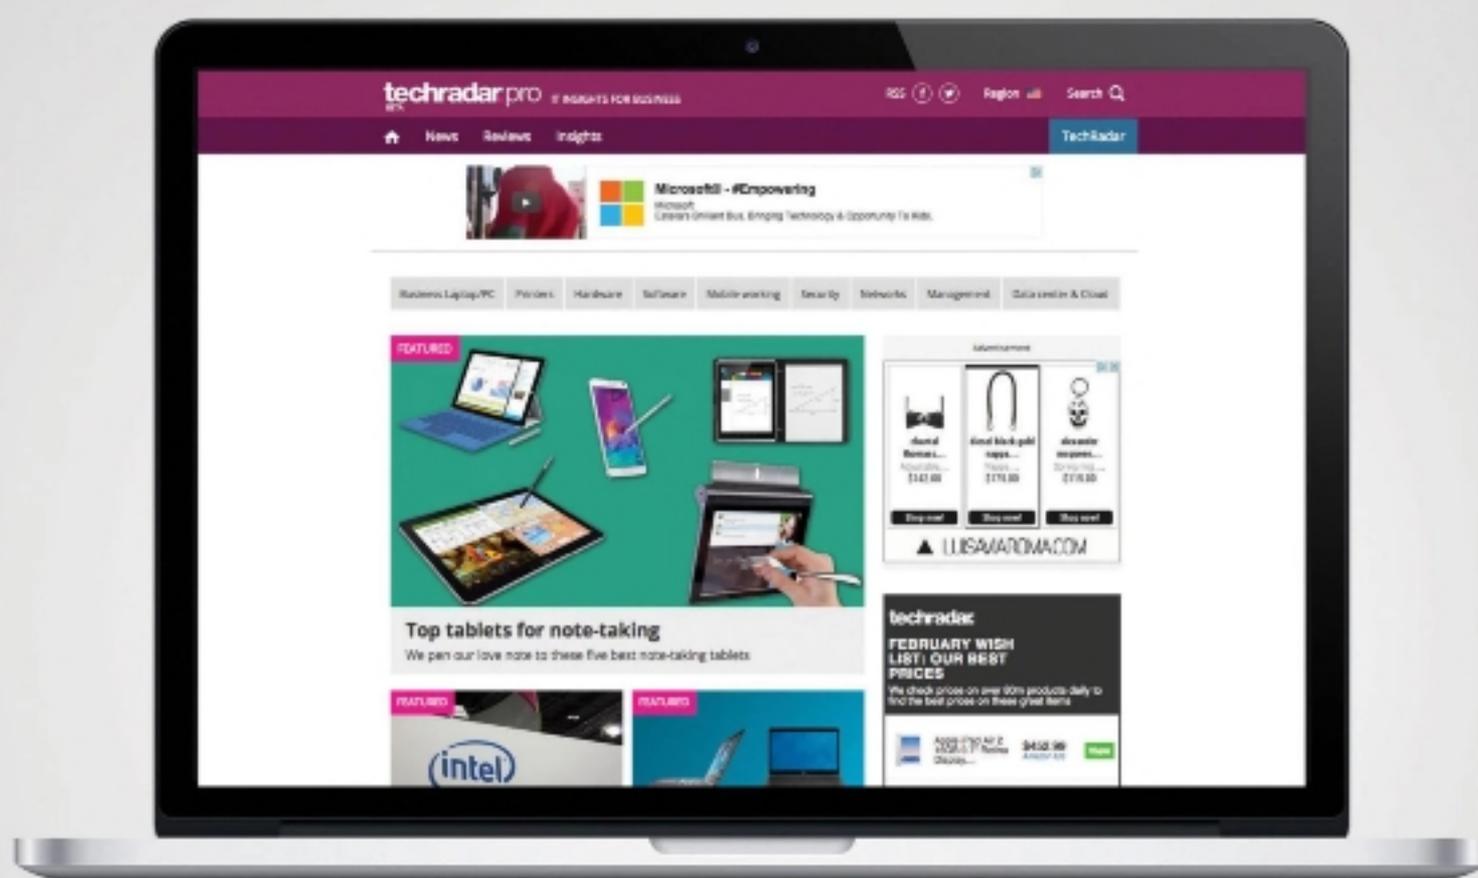
GPU	Turing TU116
Lithography	TSMC 12nm FinFET
Transistor Count	6.6 billion
CUDA Cores	1,408
Texture Units	88
ROPs	48
Core/Boost Clock	1,530/1,785MHz
Memory Capacity & Type	6GB GDDR6
Memory Speed	14Gb/s
Memory Bus	192-bit



The 1660 Super comfortably outperforms the original 1660.

# techradar.pro

IT INSIGHTS FOR BUSINESS



<http://www.techradar.com/pro>

# AMD Ryzen 5 3400G

## The new king of integrated graphics?



WITH THE 3400G, AMD might have Intel over a barrel in the graphics department. This new chip is the most sophisticated AMD APU chip on the market, with four cores supported by Radeon RX Vega 11 integrated graphics. It's also the most expensive, but not by much; the 3400G costs \$150—only \$20 more than the current price of its predecessor, the 2400G. Both use Vega 11 graphics, but the 3400G uses a 12nm architecture rather than the 2400G's 14nm. With solid memory support and an excellent 3.7GHz base clock, the 3400G looks like it might be the new king of integrated graphics.

Of course, AMD's main competitor for this chip isn't its own last-gen CPU. The opposition to Radeon Vega is Intel's UHD Graphics, with the closest comparison to Vega 11 being UHD 630, as found in CPUs such as the Core i3-8350K and i5-8600K. But Intel has a new challenger muscling in very soon: its new Iris Plus graphics, which boasts massively improved performance over UHD. Intel is shaping up to deal a serious blow to AMD's position in the APU market, so what does the 3400G do to solidify that position?

Well, those graphical cores come with 704 stream processors, and AMD's simultaneous multithreading technology means there are eight CPU threads to play with. It also gets a leg up on Intel by including an excellent stock cooler, the Wraith Stealth Spire, making the 3400G a great choice for cheap system builds. It's easy to install and compatible with 500 and 400-series AMD mobos (but some older boards will require a BIOS update).

One of the obvious downsides is that the 3400G doesn't have the 7nm Zen 2

architecture used in the higher-end Ryzen 3000 chips. Instead, it's Zen+, a sort of halfway house that's more like the original Zen format with some added extras, found in the performance models of the Ryzen 2000-series CPUs, such as the 2600X. While there's nothing wrong with the Zen+ architecture, it's now over a year old, and the lack of a Zen 2 APU makes Ryzen's lineup all the more vulnerable to the arrival of Iris Plus. Zen+ also lacks PCIe 4.0 support for superfast M.2 drives, a big disadvantage compared to the non-APU Ryzen 3000 chips.

### PRICE AND PERFORMANCE

Of course, what really matters is how well the 3400G performs in the real world. Fortunately, it's pretty good. While \$150 is a tad steep for a mid-range APU, it's decent for an eight-thread workhorse CPU, and the 3400G is good value when compared to the other 3000-series Ryzen chips. The Ryzen 7 3700X costs a little more than double the price of the 3400G, and performs twice as well on average in most tasks utilizing all cores. The 3400G does excellently in single-core tasks, still inferior to more expensive CPUs, but fantastic considering the price. Manual overclocking is unimpressive (in line with the majority of 3000-series Ryzen CPUs), but automatic overclocking using Ryzen Master will see a solid adaptive overclock of around 4GHz, squeezing an extra 5 percent or so from your performance.

Performance compared to the previous-gen Raven Ridge CPUs (in this case, the 2400G) is 5–10 percent better on average, although SSD read and write rates are relatively unchanged. This

isn't really surprising, though, because maximum transfer rates over PCIe are typically capped by the drive being used. Conversely, performance in gaming is impressive, with frame rates 20–30 percent better than the 2400G. While 1080p gaming on the 3400G isn't really viable, 720p gaming absolutely is.

Price-wise, the 3400G is 25 percent more expensive than the older 2400G—although the older APU can be found for a little cheaper than its RRP of \$130. The performance is good, yes, but in terms of raw value against power, it's not quite up to snuff. The cheaper Ryzen 3 3200G may be a better call for those looking to assemble a GPU-less system, or the incredibly cheap new Athlon 3000G. The 3400G is great, don't get us wrong, and it might be the best CPU with integrated graphics on the market right now, but products that sit at the top of their category are often not the best value, and that's the case here. —CHRISTIAN GUYTON

### VERDICT



### AMD Ryzen 5 3400G

**RYZING** Excellent integrated graphics; solid clock speeds; great stock cooler.

**FALLING** Lackluster overclocking potential; no Zen 2 architecture; not the best value APU around.

\$150, <http://amd.com>

### BENCHMARKS

	AMD Ryzen 5 3400G	AMD Ryzen 5 2400G
Blender 2.80 Beta BMW (fps)	<b>1,261</b>	1,162
Cinebench R15 Single / Multi (Index)	<b>167/885</b>	156/818
POV-Ray 3.7.1 Single/Multi (pps)	<b>403/1,934</b>	370/1,781
PCMark 8 (Index)	<b>3,690</b>	3,487
CrystalDisk 6 Sequential Read/Write (MB/s)	<b>3,492/3,277</b>	3,485/3,224
Total War: Warhammer II (fps)	<b>14</b>	10

Best scores are in bold. Our test bed consists of an ASRock X570 Taichi, 32GB of Corsair Vengeance RGB Pro DDR4-3200, a Corsair Neutron 240GB SATAIII SSD, and a Be Quiet! Dark Power Pro 1,200W. All games are tested at their highest graphical profile with AA turned on, at 1080p.

### SPECIFICATIONS

Base/Turbo Clock	3.7GHz/4.2GHz
Cores/Threads	4/8
Lithography	12nm
Cache	4MB L3
Memory Support	DDR4-2933
Memory Channels	2
Max PCIe Lanes	8
Graphics	Radeon RX Vega 11
TDP	65W



The simple black and gold finish looks great without being distracting.

# MSI MEG X570 Ace

## Ace in the hole

THE NEXT ARRIVAL in the vanguard of X570 motherboards is MSI's latest, the MEG X570 Ace. It brings all the new bells and whistles, from PCIe 4.0 to Wi-Fi 6. It's lovely looking, too, with sweet gold accents and an "infinity" RGB lighting panel that utilizes a strip of LEDs between two mirrored planes to give the appearance of an endless void reaching down into your system. It's nothing we haven't seen before, but that doesn't make this board's design any less cool.

"Cool" takes on a double meaning here, in fact, because an integrated fan, M.2 drive thermal guards, and an extended heat pipe that runs parallel to the DIMM slots mean that temperature management is top of the list for this mobo. Those M.2 thermal shields are individually detachable, too. No giant awkward panel to remove here, just a single screw and clasp, with a sticky heat pad underneath. Plus, they're really easy to remove completely if you want to use your SSD's own heatsink. This convenience shouldn't feel alien to us, yet too many X570 boards we've seen recently opt for a single large heat

plate, often connected by a cable to the motherboard itself for lighting purposes.

At \$370, it's at the steeper end of motherboard pricing, so it's got to do a lot of work to justify its existence. It sets about doing so almost immediately, though, with dual eight-pin power connectors for overclocking the CPU, and steel-shielded PCIe and DIMM slots. It certainly feels like a high-quality product, with sturdy metal parts, and very convincing lighting that we sadly can't show here (photography is a complicated business, y'all). There's a handful of questionable design decisions—why only four SATA ports, when most ATX cases have capacity for six or more?—but for the most part, it's worth the price tag.

One particularly interesting feature is the "game boost" overclocking dial. A physical eight-point dial beside the board's power and reset buttons (with a digital equivalent in the BIOS), it allows access to seven CPU overclocking presets, plus the default automatic setting. MSI bangs on a lot about power management systems in its promotional material, and for once it actually rings true; with our Ryzen 7 3700X socketed, slapping the preset to 4.3GHz got us a nice stable overclock, just below the 3700X's reported boost clock. The Ryzen 3000 chips still have problems hitting their supposed maximum clock speeds, but the Ace seems determined to squeeze all it can from the chip.

As an infomercial might say, wait, there's more! Manually overclocking the 3700X to that frequency with the Ace demands a voltage in excess of 1.3V, as

it would on most other boards. However, using the boost dial, we hit 4.3GHz across all cores with a maximum voltage of just 1.088V, even when stressing the processor with benchmarking software such as Prime96 and Cinebench. That's incredibly low for such a high frequency, resulting in stunning performance without endangering your chip with excessive voltages. The Ace demands a relatively high amount of power overall compared to some other X570 boards, so it's good to see. We used a fairly robust AIO cooler for testing purposes, but even taking that into account, we never saw any notably high core temperatures.

General performance with the Ace is solid; through all our testing, we never stumbled across any issues. Benchmarking tests had respectable results, and getting our memory running at full speed was a cinch—no meddling around with timings required. The BIOS is accessible, minus irritatingly low mouse sensitivity, and the support for Wi-Fi 6 alongside both Intel Gigabit and onboard 2.5 Gbps LAN means that connectivity is never going to be an issue. Bottom line: This is a great motherboard. It's a little on the expensive side, but if you're building a powerhouse system, this is definitely the way to go. —CHRISTIAN GUYTON

### VERDICT

# 9

### MSI MEG X570 Ace

MEGALODON Incredible auto overclocking; great power management; individual thermal shields.

MEGALOMANIA Slightly power hungry; similar performance on cheaper boards.

\$370, <http://msi.com>

### BENCHMARKS

	MSI MEG X570 ACE	ASRock X570 Taichi
Tech ARP's X264 (fps)	42.95	<b>43.3</b>
Cinebench R15 Multi (Index)	2,100	<b>2,119</b>
Fry Render (m:s)	<b>01:57</b>	02:01
AIDA64 Memory Latency (ns)	<b>75.4</b>	77.4
CrystalDisk QD32 Sequential Read (MB/s)	<b>5,001</b>	4,986
CrystalDisk QD32 Sequential Write (MB/s)	<b>4,440</b>	4,274
Power Draw Idle/Load (Watts)	60/155	<b>58/144</b>
Total War: Warhammer II (Avg fps)	<b>58</b>	54
Lowest Voltage @ 4.0GHz (Volts)	<b>1.07</b>	1.13
Maximum DC Achieved (GHz @ Volts)	<b>4.3 @ 1.09</b>	4.3 @ 1.37

Best scores are in bold. Our test bed consists of an AMD Ryzen 7 3700X, 32GB of Corsair Dominator Platinum RGB, an Nvidia GeForce GTX 1080, and a 2TB Corsair MP600 NVMe Gen4 SSD. Total War: Warhammer II was tested at the ultra preset at 1440p.

### SPECIFICATIONS

Chipset/Socket	X570/AM4
Form Factor	ATX
Memory Support	128GB @ 4,600MT/s
M.2/U.2 Support	3x M.2
SATA Support	4x SATA 6Gb/s
Max PCIe Support	3x PCIe 4.0 x16 [x8x8]
Rear I/O	2x USB 3.2 Gen1, 3x USB 3.2 Gen2, 2x USB 2.0, 1x USB 3.2 Type-C, Realtek 2.5Gb/s Ethernet, Intel Gigabit Ethernet, PS/2, Wireless A/C, 5.1 audio out, optical audio out

# TerraMaster F5-422

## Awesome NAS that won't hurt your wallet

**PROFESSIONAL-GRADE** NAS units tend to cost a pretty penny, with new features such as 10GbE support and Btrfs snapshot file security frequently driving the prices up into four figures. So, when TerraMaster revealed the F5-422, a five-bay NAS box with superfast Ethernet, extensive RAID support, and simultaneous access for up to 500 users, we were intrigued. Then suspicious, as the F5-422 costs just \$600, practically half the price of a comparable professional NAS system from industry big-names such as Synology and QNAP.

Yet, as we set up and tested this device, we were hard pushed to find any problems. Yes, the exterior is nothing special, with some cheap-feeling plastic parts, but a NAS box doesn't have to look good. The outer chassis is solid metal, sufficient to prevent significant damage from anything dropped on it. There aren't any physical security features, such as locks for the drive bays, which is a shame, but unsurprising given the price.

Setting up the F5-422 is simple. It can support up to 14TB of storage per bay (a whopping 70TB in total), be those conventional NAS HDDs or 2.5-inch SSDs—although using an SSD as a cache drive will disable one of the other drive bays. Installing the drives is easy; the brackets slot out of each bay with a front-facing clasp, and all the required screws come packaged in the box.

There's plenty in the box, in fact; a magnetic screwdriver means you don't need any additional tools, and there's

an Ethernet cable along with the power supply (which included multiple regional adapters in our version, too). There's also a pair of extra feet, should one of the unit's rubber supports go AWOL. Disassembling the F5-422 isn't too easy, but the RAM is upgradable if you're willing to go through the pain barrier.

On to the network functionality. Once the F5-422 is plugged in, there's some automated setup that can be done via the TerraMaster web front end, then you're ready to go. Sort of. While the browser app works relatively well, producing a facsimile of a desktop for you to manage your files and monitor the hardware, it has a few issues. Drag-and-drop functionality is a bit wonky, and transferring large files repeatedly in the same session causes stability issues.

Setting up direct network access through a PC's file explorer requires you to download TerraMaster's app, then directly map the network drive using the IP address. TerraMaster's online guide is easy to follow, regardless of whether you're a savvy software head. Direct access to the F5-422 via a network drive works great, with no stability problems encountered during our tests. We tested the transfer speed using the CrystalDisk benchmarking software and it hit the transfer limit of our gigabit network; there's not much danger of this machine bottlenecking your transfer rates.

The app support is incredible, too. The F5-422's Desktop TOS is compatible with

Plex, WordPress, Docker, and more, with new app connectivity coming into beta. Sync tools for Google Drive, OneDrive, and Dropbox are all in beta. It's perfect for integrating with existing cloud-based storage for small businesses or homes. Conversely, the 4K video encoding alongside apps like Plex can turn this into a great base for home entertainment (or even connecting via the mobile app for some quick movie streaming). While the F5-422 does have an HDMI port, it doesn't currently function properly, so you'll still need a system to connect to your TV.

Data safety is assured by TerraMaster, too. There are automated regular backups, multi RAID array security, AES folder encryption, and an optional cloud drive data backup. There's a wealth of customization options, which makes the F5-422 ideal for enthusiasts looking to snag their first NAS box. Yes, we'd typically recommend you build your own NAS box—custom systems are usually cheaper than pre-built—but there's just no way to preserve the compact size and effective OS; excellent features at an excellent price. —CHRISTIAN GUYTON

### VERDICT

# 9

### TerraMaster F5-422

■ **NASA** Competitively priced; numerous RAID options; fantastic app support.

■ **NASTY** No HDMI functionality; SSD caching loses bays; lack of physical security features.

\$600, <http://terra-master.com>



### SPECIFICATIONS

<b>CPU</b>	Quad-core Intel Celeron J3455 @ 1.5GHz
<b>Memory</b>	4GB
<b>Memory Capacity</b>	12GB
<b>Bays</b>	5
<b>Storage Capacity</b>	70TB
<b>Power Supply</b>	90W
<b>Ports</b>	1x HDMI, 2x USB Type-A, 2x 10GbE RJ-45, 1x 10GbE RJ-45, AC adapter
<b>Size</b>	8.9 x 8.9 x 5.4 inches
<b>Weight</b>	5.1lb

# Noctua NH-D15 Chromax Black

## A cooling colossus

**IS AIR COOLING** worth it? The folks at Noctua seem to think so, as they doggedly refuse to produce liquid coolers for CPUs, sticking with chunky metal heatsinks with fans clipped to their sides. The slick new Chromax Black model of the NH-D15 is no exception; a massive heatsink split into two towers, each connecting to six heat pipes leading to a comparatively diminutive cooling plate. From here, there are two chunky 140mm NF-A15 fans to attach, one on the side and one between the two cooling stacks.

As with many of Noctua's air coolers, the NH-D15 is sort of a hassle to set up. It's nothing too arduous, only requiring the installation of a bracket on top of the mobo's included backplate, but thick plastic washers make it a bit frustrating. It also comes with one fan pre-installed between the heatsink towers, which has to be removed to access the screwheads and attach the cooler to the motherboard. Perhaps we've become complacent, too comfortable with clip-on AMD Wraith coolers and simple AIO units.

It's surprisingly light, despite its copper heat pipes and numerous aluminum fins, but with both fans attached, it's seriously bulky. It stands six and a half inches tall, instantly excluding it from most smaller cases, and while the heatsink is shaped to allow for full use of all the DIMMs on a motherboard, the massive fans offer no such allowances. While larger boards should be able to cope with the size, standard four-DIMM ATX boards leave little room, and anything smaller is barely an option.

You can remove one fan, using only a single one between the heatsink stacks, but this does reduce the performance. The reduction in cooling is slightly less than you might expect, which is ideal, given that our test bench mobo barely has space for both. Boards with bulky rear I/O covers will have problems, forcing



the second fan to be mounted higher than it should be, potentially reducing airflow. The fans are seriously powerful, but thankfully don't make too much noise.

The box also includes a pair of low-noise adapter cables, which slightly reduce the maximum RPM to keep the noise down, and a splitter for boards without sufficient CPU fan connectors. There's a lot of stuff in the box; we love how thorough Noctua is, including not only brackets, screws, and washers for fitting the cooler to a variety of CPUs (any AMD Ryzen or Intel Core chip), but also a Phillips head screwdriver, and a tube of its excellent NT-H1 thermal compound.

Looking at the numbers, the NH-D15 is no slouch. Compared to the respectable AMD Wraith Prism (which comes with our test bench CPU, the Ryzen 7 3700X), the NH-D15 outperforms it by a moderate margin, cutting 5 C from the CPU temperature on average across a variety of tasks, both at the baseline 3.6GHz and overclocked at 4.2GHz, running at 1.1V and 1.3V respectively. The performance is admirable, but no more impressive than the average AIO cooler.

That's part of the problem here; while the NH-D15 is an excellent product, air-only cooling just can't quite measure up to AIO coolers or custom liquid loops. Noctua's cooler will run you \$140 at present, although the slightly older non-Chromax version of the NH-D15 is available for around \$100. This does somewhat raise the question of why a black paint job costs so much, but even a \$100 price point is significantly

more than plenty of good 120mm AIO coolers. In fact, it's even possible to pick up a twin-fan 240mm AIO unit from a respected manufacturer like Cooler Master or Corsair for under \$100, which is practically guaranteed to outperform most air coolers. On balance, we'd say that the NH-D15 is just about perfect if you happen to prefer air cooling, but ultimately it won't always be the best tool for the job. Don't worry though, Noctua, we still love you! —CHRISTIAN GUYTON

**VERDICT**  
**8** **Noctua NH-D15 Chromax Black**  
**CHROMAXIMUM** Excellent air cooling; works well with one fan; lots in the box.

**CHROMINIMUM** Far more expensive in black; too bulky for many cases.

\$140, <http://noctua.at>

### SPECIFICATIONS

<b>Socket Compatibility</b>	Intel LGA1150, LGA1151, LGA1155, LGA1156, LGA2011, LGA2066, and AMD AM4, AM3(+), AM2(+), FM1, FM2(+)
<b>Cooler Size</b>	140mm
<b>Cooler Dimensions</b>	6.3 x 5.9 x 5.3 inches
<b>Cold Plate Material</b>	Copper
<b>Fan Airflow</b>	82.52cfm
<b>Fan Noise</b>	26dBA
<b>Warranty</b>	Six years

# Razer Seiren Emote

## Razer gets emotional

**THIS MUST BE** one of the stranger products we've reviewed. Razer's newest microphone is essentially just the original Razer Seiren X, with one very specific addition: an eight-by-eight grid of RGB LEDs that allow for 8-bit icons to be displayed through the black mesh on the back of the microphone. It's a pretty niche feature, we're not gonna lie; a gimmick aimed at Twitch streamers who want to appear in the corner of their stream for their millions (or, maybe, tens) of screaming fans.

It actually does look pretty cool. The Emote uses Razer's new Streamer Companion App, which enables you to set up a wide variety of custom alerts using the LED display, whether that's a thumbs-up when a Twitch user subscribes to your channel, or a pulsing love heart in response to cheers. The app comes pre-loaded with dozens of animated and static emojis, even allowing for custom light responses to key phrases in the chat. It allows for a small degree of streamer-audience interaction, where the streamer can give their viewers a word to type in the chat in order to elicit a specific emoji on the microphone's display.

The degree of customization is impressive, with the app even including an editing tool for producing your own 8-bit icons. You can also upload your own 8x8 GIFs, although these can be tricky to find, and don't always behave well with the lighting effects; due to the extremely small pixel pitch, producing colors other than the standard eight RGB LED options is effectively impossible. That said, the pre-loaded icons provide more than enough variation to keep things interesting. You can also designate a "background" icon that displays on the Emote passively between alerts.

Now, to the hardware inside it. When we said that this was essentially the Razer Seiren X, we weren't being entirely truthful; while it does share the same chassis as the X—right down an identical red/green LED mute button and volume dial—the microphone inside is slightly more sophisticated. With a hypercardioid pickup pattern, it reliably picks up speech

from any position around it, which is particularly helpful, given that using it for its proper purpose will likely mean that the "front" is facing away from you.

The gain control is pretty solid, enabling the mic to pick up sound effectively when correctly set up, but it does mean that background noise sometimes breaks through. Keyboard clicking, in particular, can be quite audible in recordings using the Emote, which is quite a shame (although also a problem that the Seiren X has). Turning down the sensitivity can work to tune ambient sound out, but the bass quality suffers as a result, which can prove frustrating if you have a voice as deep as our reviewer. Belle Delphine should be safe at least, although sadly the Emote doesn't come in pastel pink.

Our main criticism isn't of the hardware, though; it's the software. Razer software has always been a mild pain—between forced installs when plugging in a new device or driver incompatibility issues, and setting up the "plug-and-play," Seiren Emote is not as straightforward as it should be. You'll need Razer Synapse 3 for microphone control settings, but you can't adjust the lighting effects using that; only the companion app can do that. A number of other elements are needlessly complicated. Setting the lighting slider to the minimum in the custom icon editor doesn't turn off that LED for some reason, instead setting it to color code 050505, which results in a dull glow. You can turn individual LEDs off properly, but you have to manually set all the color values to zero to do that.

Ultimately, this is a highly specialized product. Some people will buy into the gimmick—don't get us wrong, we actually quite like the little 8-bit display, but given that the Emote costs nearly twice as much as its predecessor, the Seiren X, we can't really recommend it wholesale. We fear that the target audience is simply too small; high-powered professional streamers will want superior audio quality, while more casual Twitch gamers might be put off by the higher price tag. —CHRISTIAN GUYTON



### Razer Seiren Emote

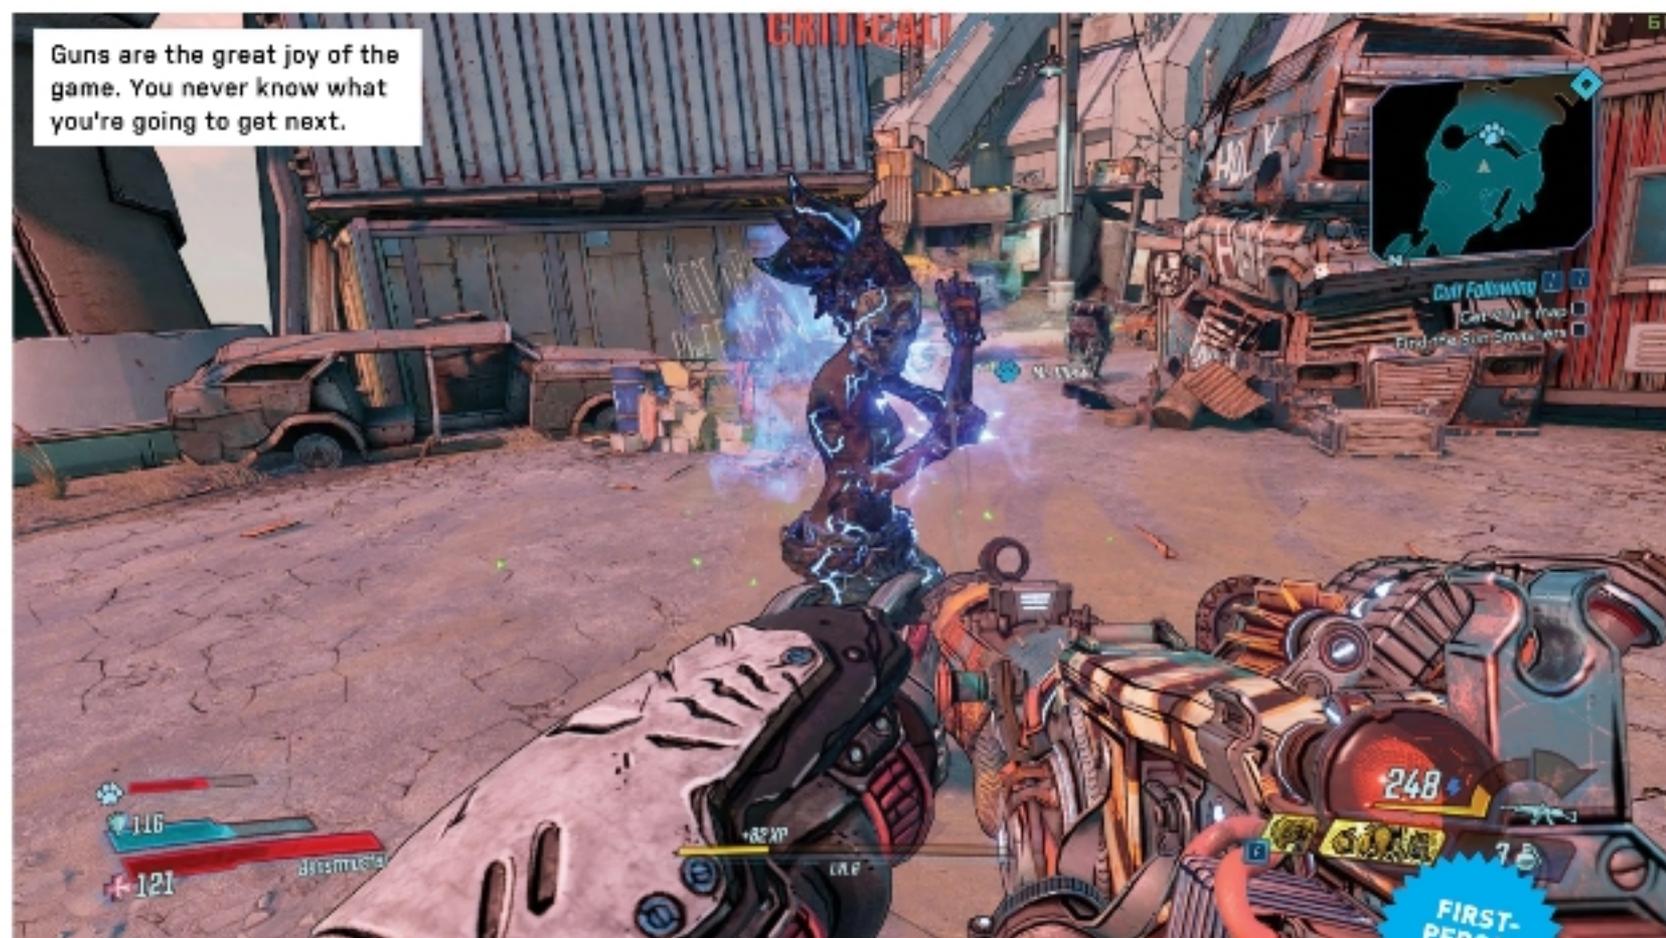
**SMILEY** Excellent audio quality; great customization options; bright display.

**FROWNY** Software is a hassle; very niche audience; not worth the extra cost.

\$180, <http://razer.com>

### SPECIFICATIONS

Frequency Response	100Hz–20kHz
Sample/Bitrate	48kHz/16-bit
Cable Length	6 feet
Polar Patterns	Hypercardioid
Element	Electret condenser
Power Consumption	5V, 500mA



# Borderlands 3

Death is just a brief setback for the financially solvent

**THERE'S SOMETHING** about twins in games at the moment. We played as them in *Wolfenstein: Youngblood*, and fought them in *Far Cry New Dawn*. In *Borderlands 3*, we just sort of sigh, wish they'd shut up, and consider skipping the cutscene.

The Calypso Twins are streamers—*Borderlands* liking nothing better than to satirize its own audience—and once they've decided to set themselves up in competition with you, you can expect a stream of obnoxious videos calling you out. *Borderlands*' world is one in which the biggest jerk rises to the top, so they're appropriate villains, but it mostly feels like they're there to keep up the barrage of knowing references to Internet life.

The game itself differs little from its predecessors. Open world, cel-shaded style, vehicles, constant weapon upgrades, and loot to be found. Missions fling you across the planet Pandora and, for the first time, beyond. Traveling to other worlds allows the series to take a break from the desert world it's been stuck with since 2009, with the jungles of Eden-6 and the stand-out metropolis of Meridian providing visual refreshment.

The rest of the game is pretty much business as usual, however. The Calypso Twins are in charge of a charming bunch

called the Children of the Vault, who provide the bulk of the cannon fodder for your Vault Hunter, who can be chosen from four initial archetypes, then customized. Vaults are caches of advanced technology, the presence of which has maintained a planet-wide economy. Not that you need to find vaults to progress—new guns, energy shields, grenades, and other loot pour from every box, corpse, washing machine, locker, and toilet.

The game is obsessed with toilet humor, but the true joy of playing it is in the weapons. The developer claims "billions" of combinations, and not knowing what you're going to pick up is the main driver behind your ransacking of every cupboard and crate. Guns that chain lightning between enemies are fun. Guns that, instead of reloading, are thrown at enemies and blow up like a grenade got us through a few encounters. The manufacturer of that last gun has a gimmick that, instead of reloading, you throw your empty gun away and a new one appears in your hand, which is possibly some sort of comment on modern life's disposable consumption, or just an excuse for a cool materialization effect.

Enemy grunts run at you, with a bit of flanking and use of cover thrown in,

but most of the time you gun them down in open areas before taking on a mini-boss to complete a mission, and getting a new gun. Fast travel is generous, as are vehicles, so getting around the large maps is no trouble.

*Borderlands 3* seems to be an attempt to out-do the previous games in all ways. It's bigger, noisier, more obsessed with the restroom, and when the jokes fall flat, they fail harder than ever. Setting up an antagonist as someone so annoying you can't wait to pop a bullet in them worked for the series before, but once the novelty and laughs wear off, you realize you're hours deep in a highly repetitive activity—especially once there's nothing left to do but collect better loot. —LAN EVENDEN

## VERDICT

# 6

### Borderlands 3

**BORDERING** Fast combat; enormous world; inventive weapons; some good jokes.

**DISORDERING** Sophomore humor; drab opening planet; Claptrap is back.

**RECOMMENDED SPECS** i7-4770/Ryzen 5 2600; 16GB RAM; GeForce GTX 1060 6GB/Radeon RX 590.

\$60, <http://borderlands.com>, ESRB: M



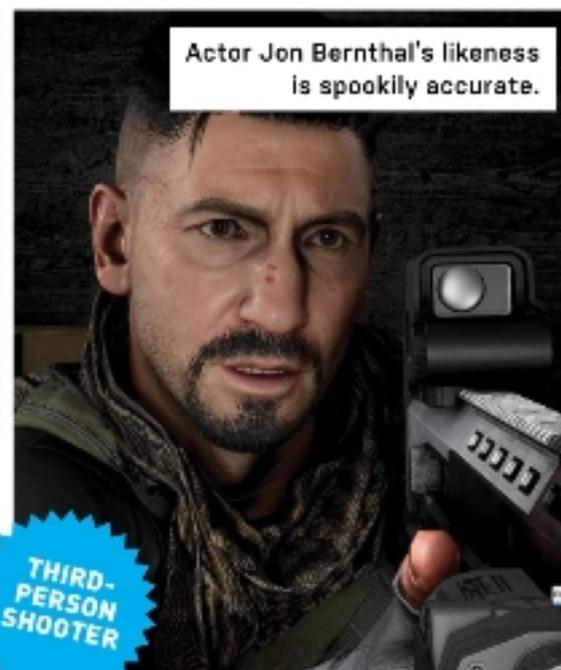
Interrogating goons gives clues to the location of enemy leaders.



Assaulting enemy bases gets you new weapons and extra currency.



Hiding by laying prone is an effective way to avoid pursuit.



Actor Jon Bernthal's likeness is spookily accurate.

THIRD-PERSON SHOOTER

# Tom Clancy's Ghost Recon: Breakpoint

## Clancy's Ghosts take on a tech bro and the Punisher

**PRETTY SOON**, we're going to need a new genre flash to describe games. Instead of "shooter" or "puzzle-platformer," perhaps we should just put "Ubisoft." Games under the Ubi umbrella trade systems and ideas like bacteria trade antibiotic resistance. This means you know exactly what you're going to get when you start up the game, whether it's *Assassin's Creed*, *Watch Dogs*, or *The Division*.

*Ghost Recon* has been part of the open-world mush since 2017's *Wildlands*. *Breakpoint* is similar in many ways, but your special forces infiltrators are in action on the fictional island of Auroa, rather than the real country of Bolivia. Much less chance of setting off a diplomatic incident that way. The bad guy this time is an ex-Ghost, Lt. Col. Walker. He and your custom Ghost Lead, call sign Nomad, have a past, and you're going to know about it whether you like it or not.

Shot down and alone, with Walker's new buddies the Wolves having taken over the island and its convenient high-tech drone factory, it's up to you to sneak around, shoot everything that moves, and

help out a bunch of farmers who are pretty handy with rifles themselves.

The game is a four-player co-op that can be soloed, but lacks the bots of previous entries in the series. If you go in alone, you're on your own, apart from a couple of drones you keep down your pants. This doesn't matter, as the game is one of stealth, maximizing the damage that can be done from one position, then moving on before you're noticed. It may take longer to clear a compound this way, but thanks to realistic gun damage, you've got a chance unless cornered and overwhelmed.

You can be a member of four classes, but they boil down to sniper or not-sniper, especially if you're playing solo. Medics may be handy for dedicated co-op players, but the flexibility of carrying a rifle and sub-machine gun, along with pistol and grenades, makes Sharpshooter the most useful class. Levels make an appearance, too, with powered-up enemies marked as especially dangerous. Equipment flies by, the island apparently groaning under the weight of guns stored upon it, with more

able to be unlocked by hunting in boxes or bought at shops and campsites.

*Ghost Recon* appears to be a natural fit for the open-world formula, but we find ourselves looking back approvingly on the days when there was a single mission per level and a hand-picked team. While you can swap out your guns at any point, power yourself up, and craft as many first aid kits, grenades, and rocket launchers as you can carry, it lacks the focus of *Recons* past, and we're not sure there's anything here that will remain long in the memory. —IAN EVENDEN

**VERDICT**  
**6** **Tom Clancy's Ghost Recon: Breakpoint**

- GHOSTLY** Good-looking; lots of collectables; enormous world to explore.
- GHOSTLY** Open world doesn't suit the franchise; bugs; lots of microtransactions.
- RECOMMENDED SPECS** Ryzen 7 1700X or Core i7-6700K; 16GB RAM; RX 5700 XT or GTX 1080.

\$60, <https://ghost-recon.ubisoft.com>, ESRB: M

# LAB NOTES

ALAN DEXTER, EXECUTIVE EDITOR



## Something Massive This Way Comes

AMD closes out 2019 with a flurry of new chips

**AMD HAS HAD A GREAT YEAR.** It's had a great couple of years, to be fair, but 2019 will go down as the year that the company really brought the fight to Intel. Zen 2, as found in the Ryzen 3000 processors, has been a huge success, and everyone, bar some die-hard gamers, is now eyeing up the red team's chips for any future build or PC purchase.

And it only looks to get worse for Intel.... At a recent presentation, AMD released the final details of the Ryzen 9 3950X, the slightly delayed 16-core, 32-thread CPU that takes on the mantle of the king of the Ryzen 3000 line. The latest addition to the family has a base clock of 3.5GHz, with a boost clock of 4.7GHz, calls on 72MB of cache, and yet maintains a relatively low TDP of 105W. AMD is recommending all-in-one cooling for this one [it doesn't ship

with an air cooler], which is a bit of shame, but understandable at this level and price point (\$749).

This wasn't the only chip talked about, though—there was also plenty of information about third-gen Threadripper, namely the 3960X and 3970X, rolling in with 24 cores/48 threads and 32 cores/64 threads respectively. The interesting snippet here is that these new Threadrippers are on a new TRX40 platform with a new sTRX4 socket, which means you'll need a new motherboard. We'll have more on these next issue.

There was yet another processor announced that caught my attention, too, and that is the Athlon 3000G. This dual-core, four-thread, 3.5GHz, unlocked chip costs a mere \$49. It may only have Vega 3 graphics, and is based on the Zen+



The Ryzen 9 3950X is imminent, but it's not the only treat AMD has lined up.

architecture, but I'm already excited about the kind of work it can be put to. That really isn't a lot of money for a fully-functioning modern CPU.



**JOANNA NELIUS**

Hardware Staff Writer

I've been spending a lot of time with Twine, a program for writing interactive fiction. I love text-based and branching narrative games, such as *Life is Strange*, so I decided to take a stab at writing my own. (I did spend eight years studying creative writing and literature.) Coding has never been my strong point, though. I've

tried learning on my own with books and online classes, but nothing has ever stuck beyond basic HTML and CSS. It's been a lot easier with Twine. Maybe it's because much of the coding in Twine is macro-based, and it's all specific to Twine. Whatever the case, I finally feel like I understand variables and other basic coding concepts.



**JARRED WALTON**

Senior Editor

Good things come to those who wait, and almost exactly one year after its PS4 and Xbox One launch, *Red Dead Redemption 2* has arrived on PC. And holy cow is it a big honking game—literally! It's eating up 112GB of my SSD. But it will all be worth it, because the PC version will be the definitive way of playing. Unfortunately, I'm currently

stuck running benchmarks instead of getting my cowboy on, and Rockstar has made my life difficult. There are no presets, the game also auto-detects your hardware and tends to adjust settings in ways you might not expect, and there are 39 individual graphics settings to play with. Definitely a case of too much of a good thing.

## Editors' Picks: Digital Discoveries

Hardware lead, Bo Moore, and staff writer, Christian Guyton, reveal what has been helping them relax



### TURTLE BEACH ELITE ATLAS AERO

Turtle Beach has been making gaming headsets for nearly 15 years, but I never paid much attention, mostly due to its tendency for over-the-top "gamer" aesthetics. However, one recent offering has caught my eye: the Elite Atlas Aero.

It's a wireless gaming headset with a few nifty features. Looks-wise, it's still got a lot of the gamer aesthetic that Turtle Beach is known for, but it's significantly less chunky, and features a nice monochrome black. Sound is good, thanks to 50mm speaker drivers and a unidirectional mic that helps only pick up your voice. It also features a software-side tweak called "superhuman hearing," which tunes game sound so that enemy footsteps, weapon reloads, and other subtle game cues are louder than normal.

Sound quality aside, the most exciting part of the Elite Atlas Aero is its earcups. Specifically, the ProSpecs system, a switch that creates a small indentation in the earcups to accommodate glasses, helping alleviate the painful pressure that comes from wearing a headset over your glasses for an extended period of time.

\$150, [www.turtlebeach.com](http://www.turtlebeach.com)



© TURTLE BEACH / BENNETT FODDY



### GETTING OVER IT

Or, to give it its full title, *Getting Over It with Bennett Foddy*. You might have heard of this game, but calling it

a game is somewhat untruthful, as I'm sure the titular Foddy would agree. The game developer's voice narrates your struggle to climb up a towering, 2D mountain using only a hammer, controlled solely by mouse movements. It's an imprecise science, and the slightest mistake can mean a severe setback as your hero launches himself merrily from the cliff face.

Foddy's narration is aware of this. He doesn't taunt you; he just punctuates your climb with observations on your progress and ruminations on the nature of game development. Sometimes, when you fail, his deadpan voice gently encourages you, reminding you that you've done this section before, and you can do it again.

*Getting Over It* is a challenge; cruel and unforgiving, often unnecessarily difficult, but never quite impossible. Almost every new obstacle will take even the savviest gamer a while to overcome, but there's no judgement when you fail. And I have. A lot. But I love it; each time I pick the game up, I get a little further, a little more convinced that I have what it takes to reach the top....

\$6, <https://foddy.net>



## Asus AiMesh AX6100 Wi-Fi System

**WE ALL NEED** Wi-Fi to survive, from checking social media to doing office work. If you're a gamer, a streamer, or you work with videos, you need it even more. And for Wi-Fi, we need routers; not very exciting, but Asus has done its best to make its latest option look cool. The AiMesh AX6100 Wi-Fi system has a symmetrical, spider-like design across two routers, designed to encompass even larger homes with a maximum coverage of 5,500 square feet.

With new high-speed Wi-Fi 6 and connectivity to any other routers in your home, the AX6100 provides total coverage, even supporting legacy Wi-Fi 5 and 4 systems. Thanks to the nifty Target Wake Time option, it can even minimize wait times for router signals for your devices, allowing them to sleep when a router signal is not required, saving battery. It also comes with Asus's AiProtection Pro, a family-oriented Internet security program that comes with all the usual bells and whistles: parental controls, blocklists, antivirus control, and more.

The AX6100 comes with tools for hardcore gamers, too. With system adaptive QoS, it can prioritize game packets over other network traffic to reduce latency. On top of that, it also uses Asus's hilariously-named WFast, software that seeks out the fastest route between your device and the game server to lower ping by up to 60 percent. It's good stuff, with an easy setup system that makes integrating the AX6100 into your existing home router system a smooth process. —CG  
\$393, <http://asus.com>

# LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

- > Windows 10 Webcams
- > Switching to AMD
- > Bargain Buys

## Webcam Problems?

A Google search suggests that the Microsoft Lifecam HD-3000 does not work with Microsoft Windows 10, either 32-bit or 64-bit. The same is also true for the Logitech C920. Were any of the eight webcams reviewed on page 19 of the October issue of *Maximum PC* tested successfully with Windows 10 64-bit?

—Robert Rue

EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS: Given that the Lifecam is a Microsoft piece of hardware, you'd assume that it would definitely work with Windows 10, but there do appear to be plenty of users who have had problems—although we never had an issue in our testing. Some reports have the webcams working fine, only to be knocked out by the big Windows updates, and then working again after others. Hardly an ideal situation. In theory, though, all of these webcams should work just fine with the built-in Windows 10 drivers, although you can use the Windows 7 drivers if necessary. As



The Logitech C920 webcam should be Win 10 compatible.

for the Logitech C920, we've had no problem with that either, although again, some people seemed to have problems when Windows 10 was first released, but this situation appears to have been fixed with an update. So, basically, while some people did have problems with their webcams with Windows 10, such issues seem to have been fixed now, and we certainly haven't encountered any problems ourselves.

## AMD Appreciation

I love the magazine—I started reading it around 2010. I've been a computer nerd since the late 1980s and 1990s, building mostly gaming PCs. During those years, I watched Cyrix,

Intel, and the more viable AMD duke it out over processors, but always stuck with my flagship, if you will: Intel. My first PC was a Packard Bell 386Sx (which was actually a huge jump from my Tandy 2000). This thinking followed me into adulthood, and my "graduated" IT career titles. It's funny, because now "nerd" isn't associated with PCs at all in the modern age, as I'm reminded as my kids fix everything mobile for me, and has taken on a new meaning, but I digress.

I was always interested in AMD, but just couldn't pull the trigger, because new builds are financially risky, more so today. I've been called many things for this—fanboy, indecisive, and/or conservative—but I honestly wouldn't chance it, even by 2017, as Ryzen rolled out. Today, however, I've happily been running a Ryzen 5 2600 for going on a year, and the earth still turns! When it came to value for money, it couldn't be beat, and deserved a serious look from the reviews as well.

I did buy Intel for all of those years, along with

Nvidia, EVGA, Asus, and so on, because of some sense of loyalty, but more because of "what worked" in my opinion. Now I look at the brave new world of equipment and have no such flagship. I can't believe AMD did it—I didn't think Ryzen/Zen would do this well, but in my experience, it's done a lot more than advertised. And it's a relief to see options—yeah, just two of them, but it's a start.

I'd like to thank you guys and gals for doing the hard work for all of us here in your publication. Word of mouth wasn't enough, and I wouldn't have broken my 20-something years of Intel loyalty without your reviews. I've written this letter, by the way, because I think there are others out there who just buy what works. That standard is vanishing now. AMD has always been as good if not better than Intel, but for me, it took this combination of price, options, and even the Ryzen format being similarly named to what I'd been accustomed to in Intel, 3/5/7, to finally change my mind.

—Ryan Anthony

↳ submit your questions to: [comments@maximumpc.com](mailto:comments@maximumpc.com)

[NOW ONLINE]

## WHY ARE GAME INSTALL SIZES GETTING SO BIG?



We recently said that traditional spinning hard drives are dead to us, but some of the latest games are trying very hard to make us reconsider. Where we used to feel that 1TB of storage was more than sufficient and 2TB was spacious, multiple games have come out or will soon arrive where the install size is more than 100GB. We certainly don't need every game we own installed on our PCs, but we do like to keep a collection around for testing and such. And with *Red Dead Redemption 2* ready to suck up 150GB, and *Call of Duty Modern Warfare* thinking it may eventually need 175GB, it got us wondering.

Just what is going on with game sizes that we're suddenly jumping from 50GB being a large install to nearly three times that much? We remember when entire worlds of gaming could be contained on just a single floppy disk—or maybe a collection of floppy disks. Whatever. The point is, if we were still doing physical distribution media, we would be looking at games that would require 20 DVDs, 250 CDs, or a mere 121,500 3.5-inch floppy disks. That's the 1.44MB "high density" variety, naturally, not the earlier 720KB version....

Read the full article: <https://bit.ly/2PPGnYu>

**EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS:** There is nothing wrong with going with what you know, but at the same time, being aware of where the great deals are is important as well. Glad to hear we've assisted you in making the switch, Ryan, and that you're happy with your Ryzen-based build. Hopefully, though, this will just result in Intel coming back stronger than ever, because competition in the market definitely benefits us all.

### Budget Builds

I have read several of your issues and I've found that none of the budget builds contain anything

cheaper than a Ryzen 5 2600—for example, a Ryzen 5 1500X or a Ryzen 3 1300X. I've only been reading *Maximum PC* at a bookstore near me since your June 2019 issue, and I'm wondering why you haven't used a Ryzen 5 1500X or a Ryzen 3 1300X in any of your builds.

—Rector-Stingray

**EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS:** We prefer to use newer hardware in our builds, because this future-proofs you the most, and even budget systems benefit from having a solid upgrade path. This is the reason we went

for the Ryzen 3 3200G in the budget build feature in the December issue, although you could drop in a previous-generation 2200G or 2400G without any changes, and save a bit of money in the process.

As for the two chips you mention, the first-generation Ryzen CPUs were impressive when released, but there have been significant improvements with Zen+ and then Zen 2 since then, so that's where we would recommend dropping your money if possible, although the first-gen Ryzens do represent great value for money. The only problem with the two chips you've specifically mentioned is that they don't have integrated graphics, so you'll need to include a discrete GPU in your spending as well, which is fine if you're looking to game, but not ideal for general use.

### Your Own Netflix

I have a customer request to build the ultimate Plex server that can rip DVDs into a ridiculous storage array to create their own "Netflix." Do you have any articles covering such a build, or would you be interested in taking on this challenge?

—Mike Cox

**EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS:** This is actually a really good idea, and one that we'd absolutely be interested in pursuing in the future. Watch this space.



The Ryzen 3 3200G is ideal for future-proofing a budget build.



VirtualBox is a great tool for running virtual machines.

### Virtual Power

I'm getting a build started and I need your help: I'm building a virtualization rig. I need to be able to run five to six virtual machines at once. I don't do much gaming, maybe a little first-person shooter (*Ghost Recon*, *Rainbow Six*, and so on), so I don't need much in the way of graphics. I want to use one of the new 3000-series AMD Ryzen CPUs. I would like to have 64GB of RAM, lots of storage, and run Ubuntu OS. I'd really appreciate any help you could give me.

—William Harrold

**EXECUTIVE EDITOR, ALAN DEXTER, RESPONDS:** Virtualization benefits from having lots of cores, plenty of memory, and a healthy amount of storage, which makes your decision to side with AMD's Ryzen chips sensible. As you have six virtual machines in your sights, we'd instinctively look toward the top of the stack, maybe even waiting for the Ryzen 9 3950X, if your budget stretches that far, although that is possibly overkill—it does depend on what you're actually doing with those VMs. Combine this powerful chip with a speedy M.2 drive, your 64GB of RAM, and a functional graphics card, and you've potentially got the core components sorted. Again, though, this would be a good topic for a future "Build It," so keep watching this space. ☺

*Holiday Season*

The perfect gift  
from just **\$7.50\***



SAVE  
UP TO  
74%

ORDER ONLINE AT

<http://bit.ly/MPCxmas197>

or call **844-779-2822** (toll free) and quote **xmas197**

\*TERMS AND CONDITIONS: Savings calculated against the full RRP (single issue price x frequency). The biggest savings are compared with overseas RRP. For all savings details, please see online. You can write to us or call us to cancel your subscription within 14 days of purchase. Payment is non-refundable after the 14 day cancellation period unless exceptional circumstances apply. Your statutory rights are not affected. All gift subscriptions will start in January 2020. Prices correct at point of print and subject to change. For full terms and conditions, please visit <http://bit.ly/maxstandc>. Offer ends December 31, 2019.

TAKE IT FROM A GEEK.™

# THE BUILDS



## BUDGET



## MID-RANGE

### INGREDIENTS

PART		PRICE
Case	BitFenix Nova TG	\$60
PSU	Corsair CX450 80+ White 450W	\$50
Mobo	ASRock B450M R4.0 AM4	\$60
CPU	AMD Ryzen 5 2600	\$120
GPU	EVGA GeForce GTX 1660 Super 6GB	<b>NEW</b> \$230
RAM	8GB (2x 4GB) Team T-Force Vulcan Z @ 3,000MT/s	<b>NEW</b> \$37
SSD	120GB HP S600 SSD	<b>NEW</b> \$24
HDD	1TB Seagate Desktop ST1000DM003	\$30
OS	Ubuntu Desktop Linux 18.04.3 LTS 64-bit	\$15

**Approximate Price:** \$626

**EVERY MONTH**, we look through the ingredients for our budget build and find certain elements that simply don't need changing. ASRock's B450M R4.0 motherboard remains one of the best-value AM4-socket motherboards on the market, and the classic 120GB Crucial BX500 is still... what? It's not on sale this month? Good grief, we'd better do something about that! It's all right; HP has a 120GB SATA III drive for just one dollar more than the BX500's original sale price, so we'll go with that. We're also swapping out the memory again, just to keep our price low. With the release of the GTX 1660 Super, it's an obvious upgrade choice; just \$10 more than our previous standard GTX 1660, while bringing 10–15 percent better performance. With a few other small savings, we've been able to significantly improve the graphical performance of this machine without raising last issue's price by a penny.

### INGREDIENTS

PART		PRICE
Case	Fractal Design Meshify C	\$90
PSU	500W EVGA 500 BR 80+ Bronze	\$50
Mobo	Gigabyte X570 UD	<b>NEW</b> \$160
CPU	AMD Ryzen 5 3600X	\$240
Cooler	AMD Wraith Spire	N/A
GPU	Gigabyte Radeon RX 5700 8GB	<b>NEW</b> \$330
RAM	16GB (2x 8GB) G.Skill Aegis @ 3,200MT/s	<b>NEW</b> \$60
SSD	500GB Samsung 970 Evo Plus M.2 PCIe SSD	<b>NEW</b> \$100
HDD	1TB Seagate Desktop ST1000DM003	\$30
OS	Windows 10 Home 64-bit OEM	\$100

**Approximate Price:** \$1,160

**WE HAD A FEW CHANGES** to make to our mid-range build this month, although most of them were quite minor. A shift to a Gigabyte X570 motherboard saved us 10 bucks, while we saved another eight dollars by making the change from G.Skill's Ripjaws V to its Aegis series. Both our case and hard drive were a little cheaper this month, too, and while the Sapphire RX 5700 rose by \$20, Gigabyte's offering stayed exactly the same price. With the cash we saved across this build, we opted to improve things somewhat by swapping out the 250GB SSD for a larger 500GB model of the Samsung 970 Evo Plus. This cost \$25 more, but our overall price for this system actually dropped by two dollars. Double the M.2 storage on our primary drive will make this build even better for gaming, with more space to store your games for improved load times via the speedy connection of PCIe 3.0.



**UPGRADE OF THE MONTH**

**MSI GEFORCE GTX 1660 SUPER VENTUS OC 6GB**



With performance that comfortably outstrips the vanilla GTX 1660, the new Super variant is a great card at a great price, ideal for new 1080p gaming builds. In fact, it sort of makes the previously appealing 1660 Ti practically pointless, with performance that falls less than 5 percent short of its big brother despite costing \$50 less. In practice, the real-world difference is almost impossible to discern in 1080p gaming, making the 1660 Super an awesome prospect. It's the best card in a crowded market right now, and while it misses out on proper ray-tracing support, it can comfortably run just about any game at above 60fps at 1080p ultra. **\$230, [www.msi.com](http://www.msi.com)**

MSI

**THE X399 PHANTOM GAMING 6** is on sale again! Like an ailing furniture store being beaten into submission by online sales, ASRock's sale price for this motherboard seems to get a bit lower every time the sale is supposed to end. With third-gen Threadrippers closing in fast, the second generation has lowered in price, too, saving us a staggering \$90 on this build from the get-go. A quick switch in cooler from 240mm to a 360mm model—more specifically, the MasterLiquid ML360 from Cooler Master—gives us improved cooling for our Threadripper 2950X CPU, and saves us nine dollars. We've also swapped out the PSU for the Corsair RM850, keeping our wattage the same as well as our price.

With those savings in place, we turned our aim to a significant upgrade elsewhere. We're making a big change, then: upgrading our GPU to the RTX 2080 Super Ventus OC from MSI. This will take our price up a fair bit, but since we've already saved over \$100 from other components, we can live with that. The change will bring this machine from a high-end 1440p system to a PC capable of proper 4K ultra gaming. It also makes ray tracing a much more viable prospect in triple-A titles, if flashy unsettlingly realistic graphics are what you're after. If that's not your bag, but this system still appeals, the RTX 2070 Super is still an awesome GPU, don't get us wrong—we just felt it was time for a change.

For more of our component recommendations, visit [www.pcgamer.com/hardware/buying-guides/](http://www.pcgamer.com/hardware/buying-guides/)

**INGREDIENTS**

PART		PRICE
Case	NZXT H700i	\$179
PSU	850W Corsair RM850 80+ GOLD <b>NEW</b>	\$120
Mobo	ASRock X399 Phantom Gaming 6	\$230
CPU	AMD Threadripper 2950X	\$680
Cooler	Cooler Master MasterLiquid ML360 RGB 360mm AIO <b>NEW</b>	\$140
GPU	MSI GeForce RTX 2080 Super Ventus OC 8GB <b>NEW</b>	\$720
RAM	32GB [2x 16GB] G.Skill Sniper X B 3,600MT/s	\$150
SSD	1TB Samsung 970 Evo M.2 NVMe	\$170
HDD	2x 3TB Seagate BarraCuda Compute	\$160
OS	Windows 10 Home 64-bit DEM	\$100

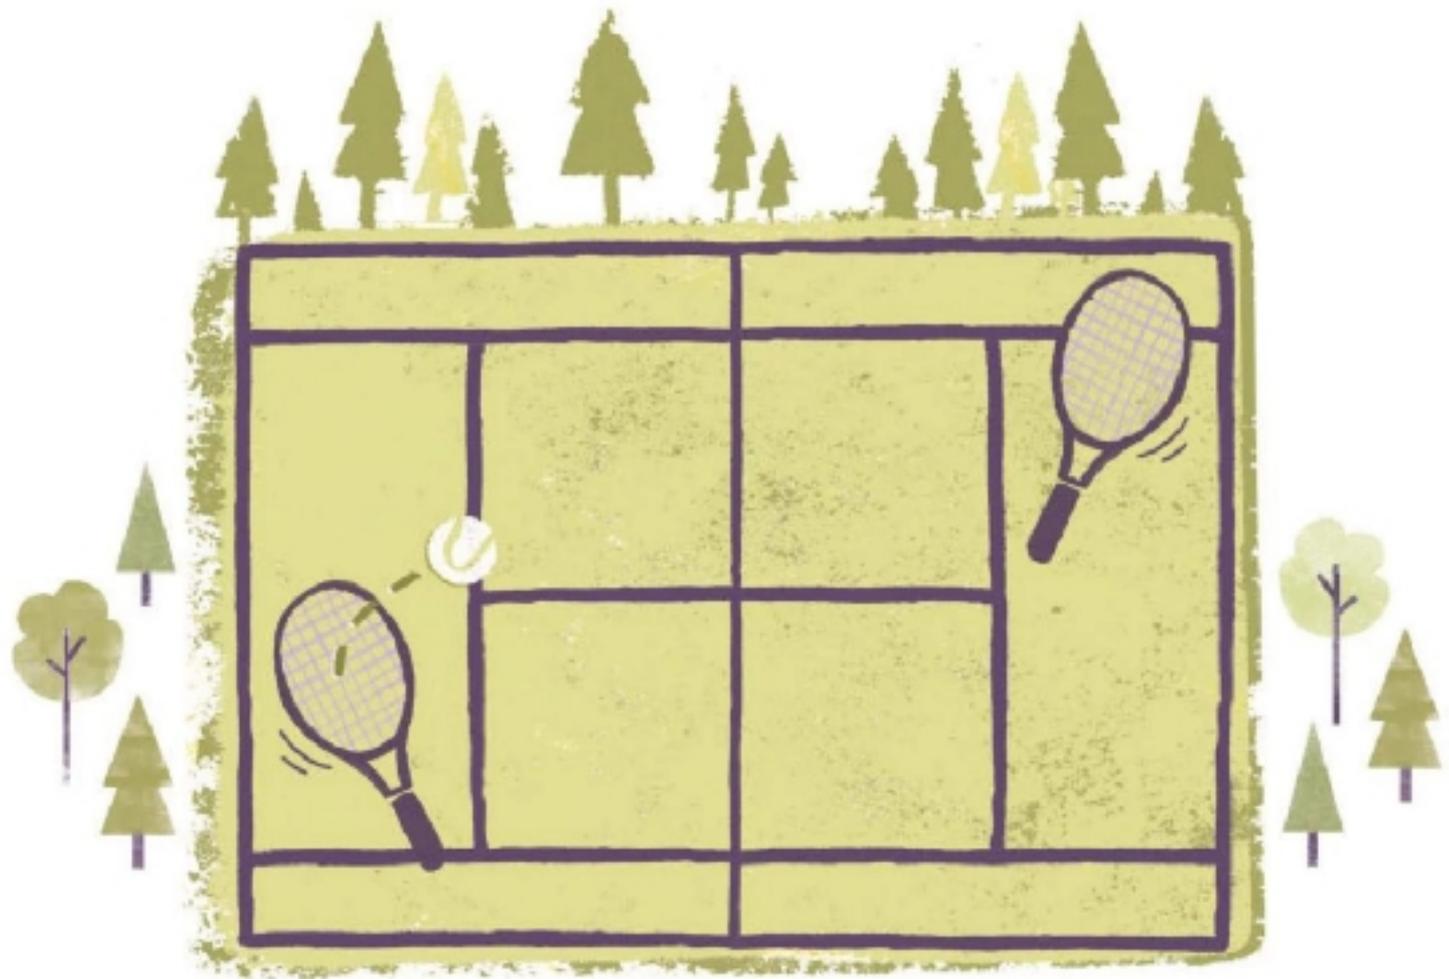
**Approximate Price: \$2,649**

Maximum PC (ISSN 1522-4279) is published 10 times a year, monthly plus a Holiday issue following the December issue, by Future US, Inc., 11 West 42nd Street, 15th Floor, New York, NY 10036, USA. Website: [www.futureus.com](http://www.futureus.com). Periodicals postage paid at New York, NY, and at additional mailing offices. Newsstand distribution is handled by Curtis Circulation Company. Basic subscription rates: one year (13 issues) US: \$24;

Canada: US\$40; Foreign: US\$40. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R123220488). PMA #40612408. Subscriptions do not include newsstand specials. POSTMASTER: Send changes of address to Maximum PC, PO Box 2024, Langhorne, PA 19047, USA. Future US, Inc. also publishes MacLife, and PC Gamer. Entire contents copyright 2019, Future US, Inc. All rights reserved.

Reproduction in whole or in part is prohibited. Future US, Inc. is not affiliated with the companies or products covered in Maximum PC. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of Maximum PC. Products named in the pages of Maximum PC are trademarks of their respective companies. PRODUCED IN THE UNITED STATES OF AMERICA.

# 22 THOUSAND TENNIS COURTS EVERY DAY!



Did you know that U.S. forests, which provide wood for making paper and many other products, have been growing every day, equal in area to more than 22,000 tennis courts!

Love magazines? You'll love them even more knowing they're made from natural, renewable and recyclable wood.

UNFAO, Global Forest Resources Assessment 2005-2015

Two Sides is a global initiative promoting the responsible use of print and paper which, when sourced from certified or sustainably managed forests, is a uniquely powerful and natural communications medium.

Discover  
more great  
#PaperFacts at  
[twosidesna.org](http://twosidesna.org)





## 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.



**LEARN MORE:** [WWW.DIGITALSTORM.COM](http://WWW.DIGITALSTORM.COM)

Digital Storm PCs Featuring Intel® Core™ i7 processors.



What would you like to do next?



Save as favorite



Search



Share

## Similar magazines



Windows 10 S:  
Microsoft's streamlined  
OS for students



Expert tutorials  
to improve  
your skills



## Other issues

**ACER CONCEPTD 500**  
Serious performance gets a hot new look **PG. 76**

**MASTER HANDBRAKE**  
Make your videos silky smooth everywhere **PG. 82**

**HISTORY OF THE PC**  
How it all began, and still defines today's rigs **PG. 48**

**OPTIMIZE YOUR SSD**  
Get the most from your superfast storage **PG. 82**

**MASTER NEXTCLOUD**  
Take control of your data with your own cloud **PG. 98**

