




## Computer Architecture

# AMD



## Dr Jeff Drobman

website



[drjeffsoftware.com/classroom.html](https://drjeffsoftware.com/classroom.html)

email



[jeffrey.drobman@csun.edu](mailto:jeffrey.drobman@csun.edu)

# AMD 10-Yr Ad 1979

Ten years ago, Advanced Micro Devices had no products, zero sales and eight of the best people in the business.

Today, Advanced Micro Devices has more than 600 products, \$200 million in sales and 8,000 of the best people in the business.

We want more. We want you.  
You'll work for the nation's fastest

growing integrated circuit company. And you'll work with people who really like to win, people who are as good at what they do as you are.

Every place has its time. Ours is now.  
Join Advanced Micro Devices.

**Catch  
the wave.**

**\$200M**

# AMD Buys \$2.1B Wafers

## News

**ADVANCED MICRO DEVICES (AMD) (154.36 +8.22) |**

### **Advanced Micro Devices to Buy \$2.1 Billion of GlobalFoundries Wafers Under Expanded Deal**

10:19 AM EST, 12/27/2021 (MT Newswires) -- GlobalFoundries (GFS) said Dec. 23 that semiconductor company Advanced Micro Devices (AMD) will now buy nearly \$2.1 billion worth of its wafers starting 2022 through 2025 under an expanded deal. A wafer is ... (MT Newswires 10:19 AM ET 12/27/2021 )

# Bob Says AMD is Good to Go

Quora



**Bob McConnell**, former VP, Embedded Processors at Advanced Micro Devices (1984-1999)



Answered Dec 3

I didn't read all the comments, but decided to answer the question. I participated in some of the Intel-AMD lawsuits. The end result of these lawsuits is that Intel is not going to be able to keep AMD from continuing to make and sell x86 processors. AMD no longer needs any information from Intel to design x86 instruction set processors. They have learned how to do it themselves. If Intel has any patents on the x86 architecture, AMD is licensed to those patents. AMD actually has a license to Intel's microcode! It's no longer relevant, but it was a big part of the lawsuit. In case you are wondering Intel tried lying, cheating, and stealing and generally behaved like scumbags.


Beyond all that, Intel is far better off with AMD getting 20% market share and providing another source. The push to ARM would be much stronger now if AMD was not in the game

# Stock Prices 2020




# Tech Titan Timeline

## Historical Perspective



# 1Q24 Revenue

---



# Sales Rankings

**2020F Top 15 Semiconductor Sales Leaders (\$M, Including Foundries)**

2020 Rank	2019 Rank	Company	Headquarters	2019 Total IC	2019 Total O-S-D	2019 Total Semi	2020F Total IC	2020F Total O-S-D	2020F Total Semi	2020/2019 % Change
1	1	Intel	U.S.	70,797	0	70,797	73,894	0	73,894	4%
2	2	Samsung	South Korea	52,486	3,223	55,709	56,899	3,583	60,482	9%
3	3	TSMC (1)	Taiwan	34,668	0	34,668	45,420	0	45,420	31%
4	4	SK Hynix	South Korea	22,578	607	23,185	25,499	971	26,470	14%
5	5	Micron	U.S.	22,405	0	22,405	21,659	0	21,659	-3%
6	7	Qualcomm (2)	U.S.	14,391	0	14,391	19,374	0	19,374	35%
7	6	Broadcom Inc. (2)	U.S.	15,521	1,722	17,243	15,362	1,704	17,066	-1%
8	10	Nvidia (2)	U.S.	10,618	0	10,618	15,884	0	15,884	50%
9	8	TI	U.S.	12,812	839	13,651	12,275	813	13,088	-4%
10	9	Infineon (3)	Europe	7,734	3,404	11,138	7,438	3,631	11,069	-1%
11	16	MediaTek (2)	Taiwan	7,972	0	7,972	10,781	0	10,781	35%
12	14	Kioxia	Japan	8,760	0	8,760	10,720	0	10,720	22%
13	15	Apple* (2)	U.S.	8,015	0	8,015	10,040	0	10,040	25%
14	11	ST	Europe	6,475	3,058	9,533	6,867	3,085	9,952	4%
15	18	AMD (2)	U.S.	6,731	0	6,731	9,519	0	9,519	41%
—	—	<b>Top-15 Total</b>		<b>301,963</b>	<b>12,853</b>	<b>314,816</b>	<b>341,631</b>	<b>13,787</b>	<b>355,418</b>	<b>13%</b>

(1) Foundry (2) Fabless (3) Includes acquired company's sales in 2019 and 2020 results.

Source: Company reports, IC Insights' Strategic Reviews database

\*Custom processors/devices for internal

AMD 4Q20 = \$3.2B





Old AMD Headquarters  
(Sunnyvale, California)

**Founder** Jerry Sanders +7  
**Headquarters** Santa Clara, California, U.S.  
**Area served** Worldwide  
**Key people** Lisa Su (President and CEO)  
 John Edward Caldwell (Chairman)  
**Number of employees** 10,100<sup>[3]</sup> (2018)

# AMD

\$6.5B

f. 1969 Tech Titan 2019

**Advanced Micro Devices, Inc. (AMD)** is an American multinational semiconductor company based in Santa Clara, California that develops computer processors and related technologies for business and consumer markets. While initially it manufactured its own processors, the company later outsourced its manufacturing, a practice known as fabless, after GlobalFoundries was spun off in 2009. AMD's main products include microprocessors, motherboard chipsets, embedded processors and graphics processors for servers, workstations, personal computers and embedded system applications.

AMD is the second-largest supplier and only significant rival to Intel in the market for x86-based microprocessors. Since acquiring ATI in 2006, AMD and its competitor Nvidia have maintained a duopoly in the discrete graphics processing unit (GPU) market.<sup>[4]</sup>



### AMD Accelerated Processing Unit



	Architecture	AMD64
<b>Cores</b>	2 to 4	
<b>Transistors</b>	32 nm 1.178b (Llano)	
	32 nm 1.303b (Trinity)	
	32 nm 1.3b (Richland)	
	28 nm 2.41b (Kaveri)	
	14 nm 4.95b (Raven Ridge)	

The **AMD Accelerated Processing Unit (APU)**, formerly known as *Fusion*, is the marketing term for a series of 64-bit microprocessors from Advanced Micro Devices (AMD), designed to act as a central processing unit (CPU) and graphics processing unit (GPU) on a single die.

AMD announced the first generation APUs, *Llano* for high-performance and *Brazos* for low-power devices in January 2011. The second generation *Trinity* for high-performance and *Brazos-2* for low-power devices were announced in June 2012. The third generation *Kaveri* for high performance devices was launched in January 2014, while *Kabini* and *Temash* for low-power devices were announced in the summer of 2013.

The Sony PlayStation 4 and Microsoft Xbox One eighth generation video game consoles both use semi-custom third generation low-power APUs.

Intel CPUs with integrated HD Graphics also have a CPU and GPU on a single die, but they do not offer HSA features (see below).



# AMD CPU's



## AMD x86

V•T•E		AMD processors	
<b>Lists</b>	Microprocessors · Microarchitectures · Chipsets · Sockets · Duron · Athlon (XP) · Athlon 64 (X2) · Sempron · Phenom · Ryzen ·		
<b>Microarchitectures</b>	<b>IA-32 (32-bit)</b>	K5 · K6 · Athlon/K7	
	<b>x86-64 desktop</b>	K8 · K9 · K10 (aka 10h) · 15h (Bulldozer · Piledriver · Steamroller · Excavator) · Zen (Zen+ · Zen 2)	
	<b>x86-64 low-power</b>	Bobcat (aka 14h) · 16h (Jaguar · Puma)	
	<b>ARM64</b>	K12 (aka 12h)	
<b>Current products</b>	<b>IA-32 (32-bit)</b>	Geode	
	<b>x86-64 (64-bit)</b>	APU · Athlon X4 · FX · Ryzen · Epyc · Opteron	
<b>Discontinued</b>	<b>Early x86 (16-bit)</b>	Am286	
	<b>IA-32 (32-bit)</b>	Am386 · Am486 · Am5x86 · K5 · K6 · K6-2 · K6-III · Duron · Athlon (XP · MP)	
	<b>x86-64 (64-bit)</b>	Sempron · Athlon 64 (X2 · II) · Phenom (II) · Turion	
	<b>Other</b>	Am9080 · <b>Am2900</b> (list) · Am29000 · Alchemy (MIPS32)	

# AMD CPU's

## AMD 386



# AMD Process Roadmap

---

At Computex, Su showed off an existing AMD desktop processor chip that had added memory stacked on top of it using the new technique. By putting the processor and the memory in closer contact, the performance of the most-demanding video games improved 15%.

"The next frontier is really to go into the third dimension, the idea of stacking chips on top of each other," Su said. The stacked chips must be carefully designed so that they don't overheat or use too much power. Comparing the process to snapping Lego bricks together but on a microscopic scale, she added: "It looks nice, but it's actually very hard to do!"

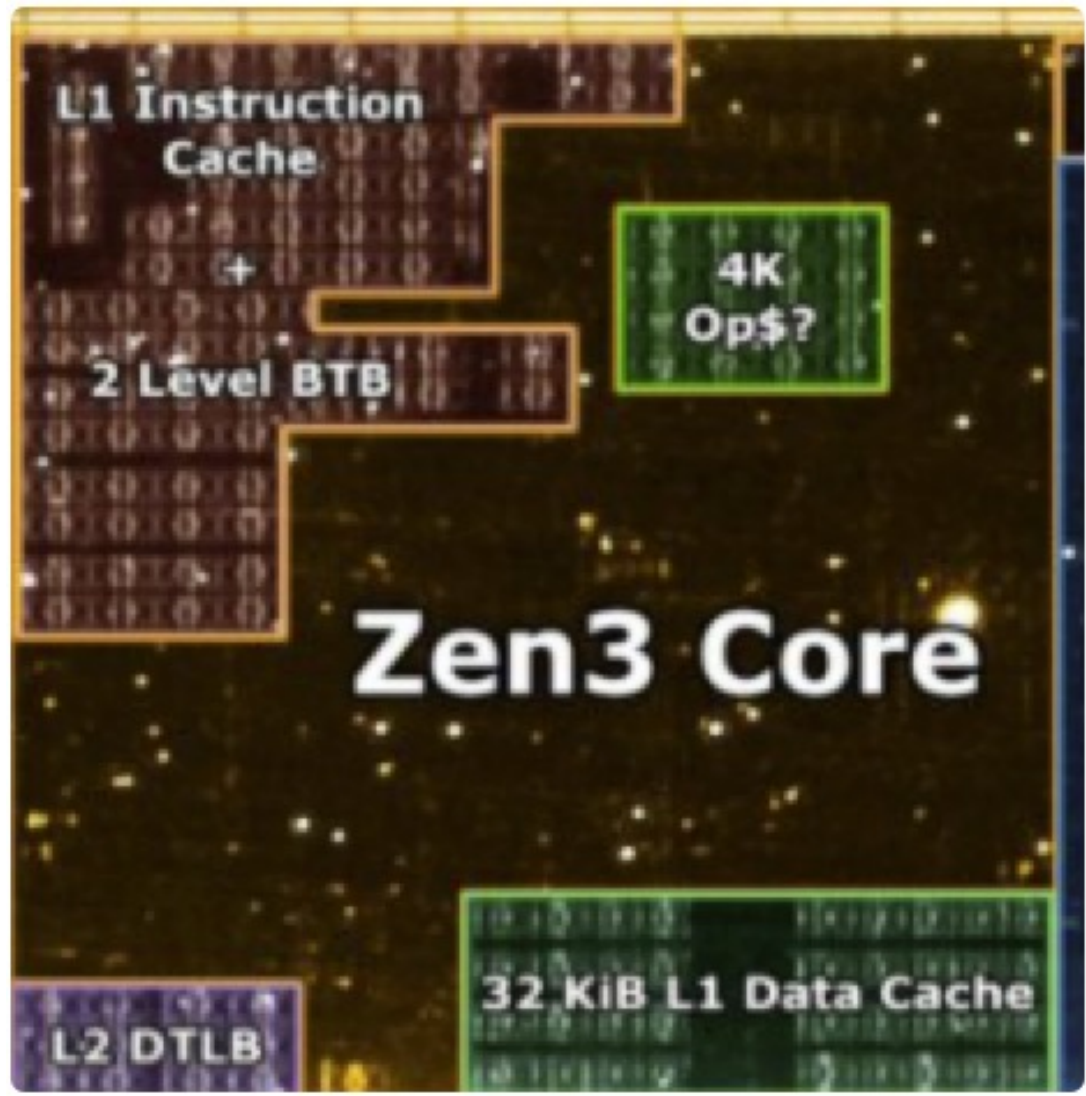
AMD will only be using the stacking technique on high-end chips, aimed at video gamers and businesses that are willing to pay extra. "This is not something that will show up for the mainstream consumer products," Su said. "It will be the top-end performance that you can get from some of our processors."

# AMD x86-64 ISA

➤ X86-64 a 64-bit version of the x86 instruction set, first released in 1999. It introduced two new modes of operation, 64-bit mode and compatibility mode, along with a new 4-level paging mode.



# AMD Zen 3 Die



# AMD Hi-End

## Threadripper



**Maheshwar S Kadi**, Technology, Rider at R-15 Bikers of Bangalore  
(2016-present)



Answered July 7

The fastest CPU in market as of now is AMD Ryzen Threadripper 3990X. below are the specs of the beast. Also, not to forget the price as of today is costing bomb of ₹ 4 Lakhs (\$ 5199).

CPU Cores: 64

Threads: 128

Max Boost Clock: Up to 4.3GHz

Base Clock: 2.9GHz

Default TDP / TDP: 280W

L1 Cache: 4MB

L2 Cache: 32MB

L3 Cache: 256MB

CMOS: TSMC 7nm FinFET

PCI Express Version: PCIe 4.0

# AMD vs. Intel

COMP122

## Which chipmaker is growing faster?

Based on these facts, it's easy to see why AMD grew at a faster clip than Intel over the past three years.

Revenue Growth (YOY)	2018	2019	2020
Intel ( <a href="#">NASDAQ:INTC</a> )	13%	2%	8%
AMD ( <a href="#">NASDAQ:AMD</a> )	23%	4%	45%

DATA SOURCE: ANNUAL REPORTS. YOY = YEAR OVER YEAR.


AMD's growth decelerated in 2019, mainly due to slower sales of gaming consoles. However, its growth accelerated again in 2020 as it launched its new Ryzen CPUs and Radeon GPUs, while Intel's ongoing chip shortage continued to generate tailwinds for both its PC and data center businesses.

Intel's overall revenue growth initially seems stable, but its data center chip sales declined year over year in the second half of 2020 -- which partly offset its higher sales of PC CPUs during the pandemic.

Intel recently outsourced the production of some of its chips to TSMC to resolve its ongoing shortages, but it doesn't expect to launch its next-gen 7nm chips until 2023. AMD launched its first 7nm CPUs back in 2019, and it will likely launch its first 5nm chips in 2022.



# AMD vs. Intel CPU's



## The key differences between Intel and AMD

Intel generated 56% of its revenue from PC-centric chips last year. It generated 36% of its revenue from data center chips, while the remaining 8% came from other types of chips -- including IoT ([Internet of Things](#)) chips, programmable chips, computer vision chips, and memory chips.

Intel agreed to sell most of its memory chip business to **SK Hynix** last October, and it's developing discrete GPUs to complement its CPUs and compete against NVIDIA and AMD. However, Intel's core business still mainly relies on its sales of x86 CPUs for PCs and data centers.

Intel manufactures its own chips with its internal foundry. But Intel's foundry fell behind **Taiwan Semiconductor Manufacturing** ([NYSE:TSM](#)), the world's largest third-party foundry, in the "process race" to create smaller chips in recent years. Those mistakes clogged up Intel's foundries and caused shortages of its latest CPUs.

AMD, which outsources the production of its CPUs and GPUs to TSMC instead of manufacturing them internally, didn't struggle with any shortages. As a result, AMD's share of the global x86 CPU market more than doubled from 18.1% to 39.4% between the first quarters of 2017 and 2021, according to PassMark Software, as Intel's market share plunged from 81.9% to 60.5%.


Last year, AMD generated nearly two-thirds of its revenue from its computing and graphics segment, which sells its Ryzen CPUs and Radeon GPUs. The rest of its revenue came from its EESC (enterprise, embedded, and semi-custom) business, which mainly sells custom chips for gaming consoles (including **Sony's** PS5 and **Microsoft's** Xbox Series consoles) and Epyc CPUs for data centers.

# AMD Epyc

CPU NAME	CORES/THREADS	BASE CLOCK	BOOST CLOCK	L3 CACHE (V-CACHE + L3 CACHE)	L2 CACHE	TDP
AMD EPYC 7773X	64/128	2.2 GHz	3.5 GHz	512 + 256 MB	32 MB	280W
AMD EPYC 7573X	32/64	2.8 GHz	3.6 GHz	512 + 256 MB	32 MB	280W
AMD EPYC 7473X	24/48	2.8 GHz	3.7 GHz	512 + 256 MB	12 MB	240W
AMD EPYC 7373X	16/32	3.05 GHz	3.8 GHz	512 + 256 MB	8 MB	240W

Meanwhile, AMD is planning to launch its Milan-X processors with 3D V-Cache technology by the end of the year. The specifications surfaced a while back. According to the source (@ExecuFix), Milan-X will consist of four SKUs ranging from 16 core all the way to 64 core. Interestingly, all of them will feature 768MB of L3 cache which is just insane, and rather excessive for the lower-end SKUs. It's possible that the final retail parts will disable part of the L3


# AMD APU's



Above is a annotated diagram of a Raven Ridge APU die, the Vega GPU and it's engines takes up well over half of the die, CCX is where AMD crams their cpu cores and caches on these hybrid chips.

# AMD Navi GPU Die

26.8B Transistors 7nm TSMC



Here's AMD's NAVI 21 GPU, which is even larger at 520 mm<sup>2</sup>, but since it's based on a 7nm process, it contains a mere 26.8 billion transistors.

# AMD in Teslas



## AMD wins first contract to supply chips to Tesla


AMD processors and graphics chips will go in the infotainment...

<https://fortune.com/2021/05/31/amd-tesla-contract-chips-infotainment-...>

The above article says on it this:

**The Tesla contract adds to AMD's momentum** following three years of gains in the computing market. **As a supplier to one of the buzziest U.S. companies**, AMD can claim significant bragging rights and yet another victory over its rivals.

The article has a lot more information on it for people who want to make a fortune! Here is another article about the same thing but with a different perspective!



## AMD confirms Tesla's new Model S and Model X will boast RDN A 2 GPUs

It appears that Elon Musk was not exaggerating when he stated in...

<https://www.teslarati.com/tesla-model-s-model-x-mcu3-specs-amd-gp...>

**With its RDNA 2 graphics system and Triple A gaming capability confirmed**, Tesla's new Model S and Model X are on track to become the vehicles with the most advanced infotainment systems on the road.


# AMD in Teslas

---

AMD processors and graphics chips will be used in the infotainment systems of newly updated Tesla Model S and Model X electric cars, which are expected to go on sale in a few weeks. With AMD's more powerful chips, Tesla owners will be able to play cutting-edge video games in their cars—when they're not driving, of course—and see maps and other items in more detail.

"The work (Tesla) is doing is truly pushing the leading edge of what you can put into a car," Su told *Fortune* in an exclusive interview before announcing the Tesla deal on Monday at the Computex computer show in Taiwan. She described the effort to make in-car dash systems as powerful as high-end PCs "part of the broader trend that computing is everywhere."

# Tesla SoC



# Section

---



# AMD History

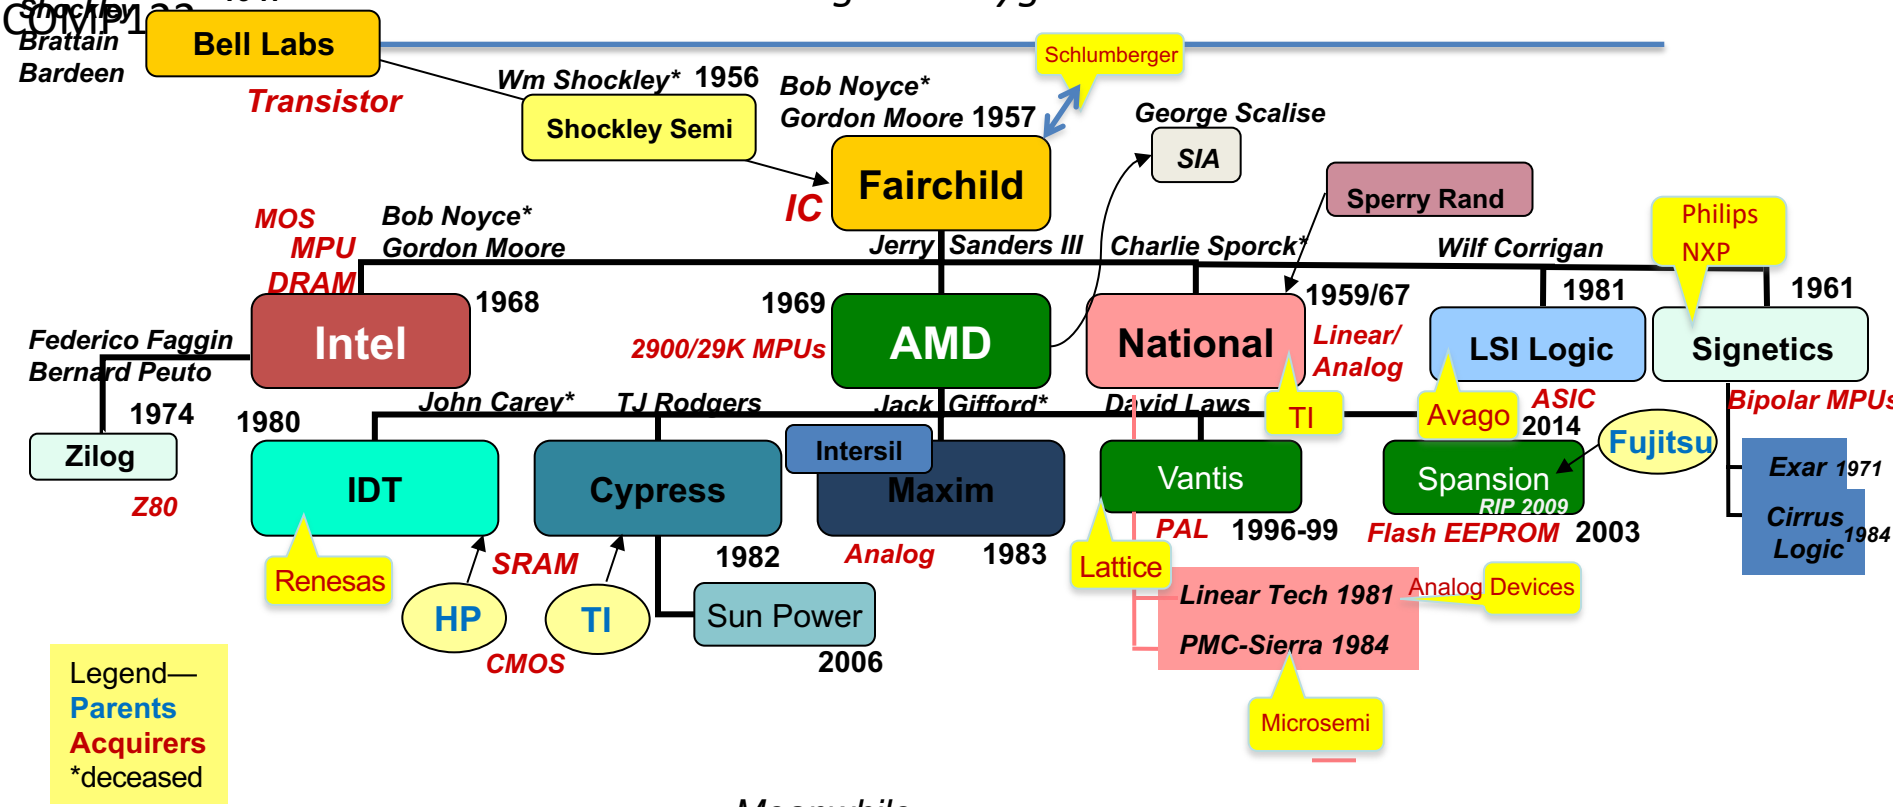
## 1969-84



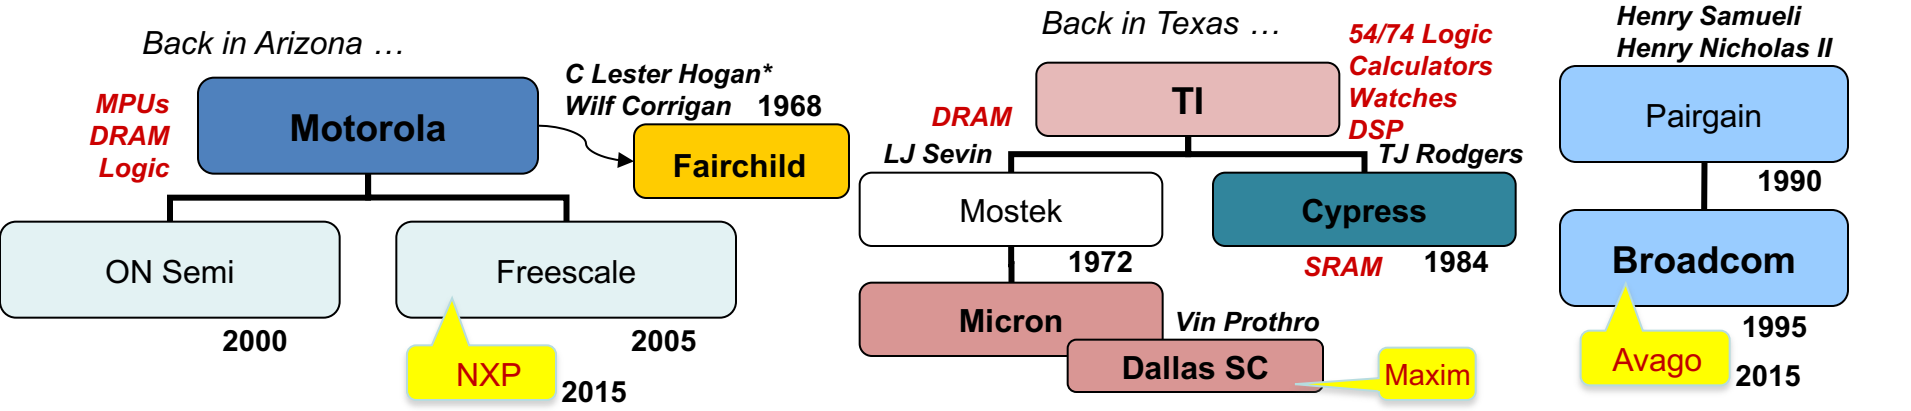
# Genesis of Silicon Valley

## The Begats & Bygones

Shockley  
Brattain  
Bardeen




### Meanwhile...



# Semiconductors & IC's

## MILESTONES

Year	Event	Process
1968	❖ <b>Intel</b> founded	
1969	❖ <b>AMD</b> founded	
MPU 1971	❖ <b>Microprocessor &amp; RAM</b> in <b>MOS</b> invented by Intel (i4004/i8008, i1101/3)	
1972	❖ Ion Implantation (replaces chemical diffusion)	Mfg -- process
1974	❖ Digital <b>Watch</b> IC invented	
1978	❖ <b>Wafer stepper</b> invented for fabs	Mfg -- process
1979	❖ <b>IDT</b> founded ( <b>CMOS</b> )	Mfg -- process
ASIC 1981	❖ <b>LSI Logic</b> founded ( <b>ASIC</b> )	
FPGA 1984	❖ <b>Xilinx</b> makes 1 <sup>st</sup> <b>FPGA</b> , <b>MIPS</b> founded as early <b>RISC</b> pioneer (licenses LSI & IDT)	
1985	❖ <b>ARM</b> founded as Acorn RISC Machines	
Flash 1987	❖ <b>Toshiba</b> intro's <b>Flash</b> EEPROM, <b>TSMC</b> founded (foundry)	Mfg -- process
GPU 1998	❖ <b>Nvidia</b> founded (1993) – 1 <sup>st</sup> <b>GPU's</b> (1998)	
2002	❖ Intel goes to <b>300mm</b> (12in) wafers	Mfg -- process
2009	❖ AMD spins off fabs to <b>Global Foundries</b> (owner Abu Dhabi)	Mfg -- process
2011	❖ Intel <b>FinFET</b>	Mfg -- process
Dec 2019	❖ Intel intro's 1 <sup>st</sup> <b>QC</b> chip ("Horse Ridge")	
2020	❖ ARM intro's "backside power" process	Mfg -- process



# AMD



1969





From left: W. Jerry Sanders III, President and Chairman of the Board. D. John Carey, Managing Director of Complex Digital Operations. Sven E. Simonsen, Director of Engineering, Complex Digital Operations. Frank T. Botte, Director of Development, Analog Operations. James N. Giles, Director of Engineering, Analog Operations. Edwin J. Turney, Director of Sales and Administration. Jack F. Gifford, Director of Marketing and Business Development. R. Lawrence Stonger, Managing Director, Analog Operations.



<p><b>Chuck Keough</b> mid-america area sales manager Chicago (312) 297-4115/6</p>	<p><b>Steve Marks</b> eastern area sales manager New York (212) 343-2220/1</p>	<p><b>Ed Turney</b> director of sales coordinates all sales activities Sunnyvale</p>	<p><b>Steve Zelencik</b> western area sales manager Los Angeles (213) 360-2102/3</p>	
<p><b>Len Brown</b> linear product marketing LIC pricing, specs, product plans &amp; introductions Sunnyvale</p>	<p><b>Elliott Sopkin</b> communications ad, magazine reprints &amp; all printed materials Sunnyvale</p>	<p><b>Jerry Sanders</b> president coordination &amp; implementation of AMD goals/objectives Sunnyvale</p>	<p><b>Jolene Trout</b> customer service, delivery, scheduling Sunnyvale</p>	<p><b>Shel Schumaker</b> digital marketing/headquarter sales, DIC pricing, specs, new product, coordination of distributor &amp; international activities Sunnyvale</p>

Bell Labs

# Founders HoF



Wm Shockley



Fairchild founders (8)



Wilf Corrigan

Fairchild  
 Chairman/CEO,  
 LSI Logic  
 founder




AMD co-  
 founder



Jack Gifford

In 1983, Gifford co-  
 founded [Maxim Integrated  
 Products](#)



Jerry Sanders  
 CEO, AMD  
 1969-2002

From left: W. Jerry Sanders III, President and Chairman of the Board. D. John Carey, Managing Director of Complex Digital Operations. Sven E. Simonsen, Director of Engineering, Complex Digital Operations. Frank T. Battie, Director of Development, Analog Operations. James N. Giles, Director of Engineering, Analog Operations. Edwin J. Turney, Director of Sales and Administration. Jack F. Gifford, Director of Marketing and Business Development. R. Lawrence Stonger, Managing Director, Analog Operations.



Bob Noyce



Gordon Moore



T. J. Rodgers

Cypress Semi founder

# AMD

TECH HISTORY ARTICLE BY DR JEFF DROBMAN



### ❖ AMD Founders

The 8 AMD founders, all Fairchild engineers, are listed here:

1. **Jerry Sanders** – *Chairman, President and CEO* (remained in those positions until the twilight; retired in 2002 at age 65)
2. **D. John Carey** – *Dir. Digital Operations* (left to be CEO of Integrated Device Technology ca. 1980)
3. **Edwin J. Turney** - *Dir Sales and Admin* (left in 1974, to be replaced by "Z") [b. 1929– d. 2008]
4. **Sven-Erik Simonsen** – *Dir Engineering, Digital Operations* (remained for many years) [Ed. note: my direct boss]
5. **John "Jack" Gifford** *Dir Mktg and Bus Dev* (left to be CEO of Maxim Integrated Products; BS UCLA 1963, MS UC Berkeley [b. 1941–d. 2009])
6. **Jim Giles** *Dir Engineering-Analog Operations*
7. **Frank T. Botte** *Dir Dev-Analog Operations*
8. **Larry Stenger** *Managing Dir-Analog Operations*



### WJ "Jerry" Sanders III (b. 9-12-1936) – *Chairman, President and CEO*

As stated in the ad, Jerry spent his first 8 years in semiconductor IC sales and marketing, after graduating magna cum laude from the University of Illinois in electrical engineering. Jerry's first job was a brief stint as a design engineer for Douglas Aircraft in Los Angeles. Jerry early on got the "sales" itch, and went to work as a Sales Engineer for Fairchild (ca. 1960) in the Los Angeles area. Jerry got promoted to District Sales Manager and opened an office in Hollywood. There, he worked with Ed Turney, his first Sales Engineer. After the Hogan

# AMD & IC History 1969-84



First 15 Years

A Case Study In Realizing The American Dream:

Sanders and Advanced Micro Devices:  
The First Fifteen Years, 1969--1984

Copyright (c) 1984, 2004 by Thomas A. Skornia  
All Rights reserved under International and  
Pan American Copyright Conventions

# AMD & IC History 1969-84



First 15 Years

## ABOUT THE AUTHOR

Tom Skornia served as legal counsel at the initial organization of AMD in 1969. He continued in that capacity until joining the company full time as Vice President and General Counsel from 1977 through 1982. In these roles, he was involved in or privy to every major event at AMD for thirteen of its first fifteen years. For the next two years, he followed its developments closely from outside while he prepared the original of the manuscript for this book.

Skornia is an alumnus of Grinnell College, as was the late Bob Noyce,

# AMD & IC History 1969-84



1968

First 15 Years

The State of Fairchild

The Mot invasion: **Hogan's Heroes**

## HOGAN'S HEROES

In the summer of 1968, Sherman Fairchild took stock of his namesake company and found it wanting. After a decade, and without impressive competition in its chosen field, the company still had a poor reputation for on time delivery and product reliability. Its strong suits were marketing and sales and technological innovation. Jerry Sanders' hard-charging sales guys had contributed greatly to a massive broadening of the applications for transistors. Fairchild scientists had created the seminal planar process. Founder Bob Noyce had invented the integrated circuit itself. The honor of this latter breakthrough, however, ultimately had to be shared with Texas Instruments' Jack Kilby to settle a lawsuit brought by TI over the question of credit for that invention.



# AMD & IC History 1969-84

1968

First 15 Years

The Mot invasion: **Hogan's Heroes**

Dr C. Lester Hogan → Wilf Corrigan

Fairchild opted for a group he identified as more manufacturing managers than engineers to address the delivery and reliability problems. Still, the seven top managers he brought in from Motorola were themselves engineers. And their leader, C. Lester Hogan, was a fairly prominent semiconductor technologist in his own right. Interestingly, two of these "Hogan's Heroes" later became high level managers at AMD. They joined in 1974, roughly concomitant with a palace coup by another Hero, Wilfred J. Corrigan, who succeeded in displacing Hogan himself as President and C.E.O. of Fairchild.

The "Hogan's Heroes" label was actually adopted by a group of Fairchild sales guys for themselves upon their hearing of the intended Hogan takeover while they were attending the Fairchild annual sales conference in Hawaii in

# AMD & IC History 1969-84

1968

First 15 Years

Exit Intel Founders 3: Aug 1968

THE INTEL TROIKA:

Noyce, Moore, Grove

Sherman Fairchild's mid-game management substitution set off an expectable reaction. Bob Noyce, Gordon Moore and Andy Grove (the top two technologists and the top manufacturing guy) left Fairchild to form Intel Corporation in August 1968. They were joined over succeeding months by many other Fairchild supporting stars. This was neither the first nor the last important Fairchild semiconductor spinoff. But it was a watershed in that process in eviscerating Fairchild's technology staff and forever dooming its efforts to enter the MOS business as an effective competitor. (An earlier spin off to National Semiconductor of Charlie Sporck, followed by Fred Bialek, Pierre Lamond, Ed Pausa and Don Valentine, similarly set back Fairchild's ability to manufacture effectively at low cost.) Take away technology and low cost manufacture, then chase out the marketing and sales

# AMD & IC History 1969-84

1969

First 15 Years

## IN THE BEGINNING

FEBRUARY 21, 1969:

The intercom buzzed.

"There's a guy on the phone who says he was referred by Bob Noyce.

Needs a business lawyer. Want to talk to him?"

"This is Jerry Sanders. A group of us are starting a company and we're looking for a lawyer who knows something about how to do that to help us with it. Are you interested?"

"Absolutely. When would you like to come in and talk about it?"

"Three of us, John Carey, Ed Turney and I, are coming up to the City on Monday to interview some candidates for our attorney. When do you have time?"

"How about lunch, and then we can come back to the office to continue, if we need to?"

# AMD & IC History 1969-84



AMD Founding

First 15 Years

MALIBU

Dec 1968: Malibu Colony Beach House

(I know, I was there for my UCLA Xmas break.)

For the now somewhat depressing 1968 winter holidays, Sanders withdrew to one of his favorite spots in the world. In Malibu, in Southern California, he settled in to contemplate his situation and his future. The setting: a beach house with a \$600 monthly rental tab, a princely sum for an unemployed sales type. Crucial momentum had been lost, and now at age 32, with three different jobs already behind him (McDonnell Douglas, Motorola and Fairchild) he needed a nimble and promising option just to stay even, much less to advance by this adversity. It was to be a recurrent formula for Sanders from that time on.

How better than to become president of a company? The somewhat annoying problem was that no presidency of an attractive company was open. But what shortly appeared to be open was the pathway to forming an entirely new company, whose presidency Sanders could occupy from the beginning.

# AMD & IC History 1969-84



AMD Founders 8

First 15 Years

ENTER THE KING AND COURT

Two others, who went over the side at Fairchild shortly after Sanders, were an English production type named John Carey, and an old Sanders sales sidekick, Ed Turney. A phone call from Malibu to Lake Arrowhead, where Turney was formulating a plan to enter the real estate business, opened a dialogue on the subject of starting a company. Together, Sanders, Turney and Carey proposed to recruit a technologist to round out a strong digital bipolar semiconductor team. Their first choice, Clive Ghest, turned them down. Their second choice, a Danish-national technologist still at Fairchild, Sven Simonsen, agreed only after the application of a large measure of the Sanders' charm and persuasion. (Somewhat paradoxically, after Simonsen's departure from AMD in October, 1982, Ghest, who had himself joined AMD in 1971, left and then rejoined in 1975, took over Sven's title as Vice President and Technical Director.)

These four, it turned out, would be "a little more equal" than the other four founders, and would emerge as much more influential in the

# AMD & IC History 1969-84



First 15 Years

Threat from Japan: **DRAM**

Generation	Market Share
4K	5%
16K	40%
64K	70%

Silicon Valley Response:

**Semiconductor Industry Assn**

# AMD & IC History: SIA

First 15 Years

Semiconductor Industry Assn

On the evening of March 16, 1977, in a private room in Ming's Chinese Restaurant in Palo Alto, California, these four met over dinner, together with Sanders' counsel, to plot strategy and form an organization to respond to the Japanese challenge. Within a week, incorporation papers were filed in Sacramento for the first ever, dedicated semiconductor trade group, the Semiconductor Industry Association. Included among the Directors on the incorporation papers, in addition to Noyce, Corrigan, Sanders and Sporck, was Motorola's John Welty, avidly recruited by Noyce to provide big company

# AMD & IC History: SIA

First 15 Years

## THE CRUDE OIL OF THE '80'S

For the longer term, it was clear that technological as well as political cooperation among U.S. manufacturers would be required. Its beginnings were already present in SIA's research arm, Semiconductor Research Corporation, founded in 1982, and the even more aggressive Microelectronics and Computer Corporation, initiated by Control Data Corporation's hard-charging CEO, Bill Norris. A key player respecting these initiatives was the U.S. antitrust establishment, whose first response under the Carter Administration was hostile. Thankfully, there later evolved a much more cooperative government attitude which was extended to support for 1984 legislation greatly liberalizing the rules for the safe conduct of joint basic research.



# AMD & IC History: SIA

## SIA in 1984

From these beginnings, SIA went on by 1984 to encompass more than 50 U.S. manufacturers and to develop and publish periodic international statistical reports on the semiconductor industry which include the Europeans, the Japanese themselves, and even Texas Instruments, which was one of the last holdouts. For a geographically concentrated industry with a relatively small employment base and no labor unions, the SIA established remarkable influence and credibility in Washington. Much of its program of increased research and development credits, resistance to tougher export controls and a form of copyright-like protection on semiconductor chip topography had been implemented, and there was much more to come.

The jury had not even retired to deliberate yet on the question of comparative effectiveness of Japanese versus U.S. semiconductor manufacturers, but it was clear that in this case for the first time, the Japanese were up against a target which was responding in a timely and effective fashion to their challenge.


# AMD vs. Intel Legal Battles

## Litigation with Intel [\[ edit \]](#)

See also: *Intel Corp. v. Advanced Micro Devices, Inc.* and *Advanced Micro Devices, Inc. v. Intel Corp.*

AMD has a long history of litigation with former (and current) partner and x86 creator [Intel](#).<sup>[266][267][268]</sup>

- In 1986, Intel broke an agreement it had with AMD to allow them to produce Intel's micro-chips for [IBM](#); AMD filed for [arbitration](#) in 1987 and the arbitrator decided in AMD's favor in 1992. Intel disputed this, and the case ended up in the [Supreme Court of California](#). In 1994, that court upheld the arbitrator's decision and awarded damages for breach of contract.
- In 1990, Intel brought a copyright infringement action alleging illegal use of its 287 microcode. The case ended in 1994 with a jury finding for AMD and its right to use Intel's microcode in its microprocessors through the 486 generation.
- In 1997, Intel filed suit against AMD and [Cyrrix Corp.](#) for misuse of the term [MMX](#). AMD and Intel settled, with AMD acknowledging MMX as a trademark owned by Intel, and with Intel granting AMD rights to market the AMD K6 MMX processor.
- In 2005, following an investigation, the Japan Federal Trade Commission found Intel guilty on a number of violations. On June 27, 2005, AMD won an [antitrust](#) suit against Intel in Japan, and on the same day, AMD filed a broad antitrust complaint against Intel in the U.S. Federal District Court in [Delaware](#). The complaint alleges systematic use of secret rebates, special discounts, threats, and other means used by Intel to lock AMD processors out of the global market. Since the start of this action, the court has issued [subpoenas](#) to major computer manufacturers including [Acer](#), [Dell](#), [Lenovo](#), [HP](#) and [Toshiba](#).
- In November 2009, Intel agreed to pay AMD \$1.25bn and renew a five-year patent cross-licensing agreement as part of a deal to settle all outstanding legal disputes between them.<sup>[269]</sup>





# Section

---



# Products



# AMD Product Lines

## Advancing Enterprise AI

AI Platforms from Cloud to Edge to Endpoints

### AMD INSTINCT

Advance Generative AI with AMD Instinct™ Accelerators

Explore how AMD Instinct™ accelerators can supercharge large, dedicated AI deployments.

### AMD EPYC



Advance Data Center AI with AMD EPYC™ Processors

Harness the power of AI to deliver business insights with servers powered by AMD EPYC™ processors.

### AMD RYZEN AI

Advance Your Business with AMD Ryzen™ AI

Experience the power of AI PCs, and be ready to usher in the future with new capabilities.





# Processor Products



June 2024

<b>AMD</b>	<b>Products</b>	<b>Solutions</b>	<b>Resources &amp; Support</b>	<b>Shop</b>
<b>Processors</b> >	<b>Servers</b> EPYC	<b>Business Systems</b> Laptops Desktops Ryzen AI for Business	<b>Workstations</b> Ryzen Threadripper Ryzen PRO for Mobile Workstations Ryzen	
<b>Accelerators</b>				
<b>Graphics</b>	<b>Personal Laptops</b> AMD Advantage Premium Ryzen with Radeon Graphics Athlon with Radeon Graphics Ryzen AI for Consumer	<b>Personal Desktops</b> AMD Advantage Premium Ryzen Athlon with Radeon Graphics	<b>Handheld</b> Ryzen Z1 Series	
<b>Adaptive SoCs, FPGAs, &amp; SOMs</b>				
<b>Software, Tools, &amp; Apps</b>				

**Embedded**  
EPYC and RYZEN  
Partner Ecosystem  
Industry Solutions

**Processors**

**Accelerators** >

**GPU Accelerators**  
Instinct Accelerators  
Documentation



# Graphics Products



June 2024

## Workstations

Radeon PRO

## Desktops

AMD Advantage Premium

Radeon RX

## Laptops

AMD Advantage Premium

Radeon Mobile Graphics

Processors

Accelerators



## GPU Accelerators

Instinct Accelerators

Documentation



# AMD Products

June 2024

Model	Cores / Threads	Boost <sup>9</sup> / Base Frequency	Total Cache	PCIe®	TDP
AMD Ryzen™ 9 9950X	16 / 32	Up to 5.7 GHz / 4.3 GHz	80MB	Gen 5	170W
AMD Ryzen™ 9 9900X	12 / 24	Up to 5.6 GHz / 4.4 GHz	76MB	Gen 5	120W
AMD Ryzen™ 7 9700X	8 / 16	Up to 5.5 GHz / 3.8 GHz	40MB	Gen 5	65W
AMD Ryzen™ 5 9600X	6 / 12	Up to 5.4 GHz / 3.9 GHz	38MB	Gen 5	65W

Model	Cores / Threads	Boost <sup>9</sup> / Base Frequency	Total Cache	PCIe®	TDP
AMD Ryzen™ 9 5900XT	16 / 32	Up to 4.8 GHz / 3.3 GHz	72MB	Gen 4	105W
AMD Ryzen™ 7 5800XT	8 / 16	Up to 4.8 GHz / 3.8 GHz	36MB	Gen 4	105W

# AMD AI Products



Model	Cores / Threads	Boost <sup>9</sup> / Base Frequency	Total Cache	Graphics Model	cTDP	NPU
AMD Ryzen™ AI 9 HX 370	12 / 24	5.1 GHz / 2.0 GHz	36MB	AMD Radeon™ 890M Graphics	15-54W	Yes (50 TOPs)
AMD Ryzen™ AI 9 365	10 / 20	5.0 GHz / 2.0GHz	34MB	AMD Radeon™ 880M Graphics	15-54W	Yes (50 TOPs)

Model	Compute Units	AI Accelerators	Ray Accelerators	Memory	TBP (Watts)	SEP (USD)
AMD Radeon™ PRO W7900 Dual Slot	96	192	96	48GB GDDR6 ECC	295W	\$3,499



# Products



Ads



AMD Ryzen 9  
3900X 3.8 Ghz  
12-Core AM4...  
**\$489.99**  
✓ B&H Photo-...  
★★★★★ 461



AMD Ryzen 9  
3950X 3.5 Ghz  
16-Core AM4...  
**\$723.99**  
✓ B&H Photo-...  
★★★★★ 100



Amd Ryzen 9  
3950X Hexadeca  
Core 16 Core...  
**\$802.13**  
✓ PCNation  
★★★★★ 100



Amd Ryzen 9  
3900x 12-Core,  
24-Thread 4.6...  
**\$489.99**  
Walmart  
★★★★★ 9



Amd Ryzen 9  
3900Xt 12-Core,  
24-Thread...  
**\$649.00**  
Walmart  
★★★★★ 3



Amd Ryzen 9  
3900Xt 12-Core,  
24-Thread...  
**\$599.00**  
Walmart  
★★★★★ 3



Amd Ryzen 9  
3950X 16-Core,  
32-Thread 4.7...  
**\$739.00**  
Walmart  
★★★★★ 100



AMD Ryzer  
3950X 3.5Ghz  
16-Core AM4...  
**\$749.00**  
✓ Adorama Ca...  
★★★★★ 100



**Extreme Performance for Gamers, Streamers and Creators**



**Introducing GameCache on 3rd Gen AMD Ryzen™ Processors**



## Videos of Ryzen 9

[bing.com/videos](https://bing.com/videos)

The video player shows a thumbnail for 'AMD Ryzen 9 5900X Benchmark Review'. The thumbnail features a large 'RIP' text over an Intel i9-10900K processor, with an AMD Ryzen 9 5900X processor shown next to it. The video title is 'AMD Ryzen 9 5900X Benchmark Review', it has 113K views, and was posted 3 months ago. The channel is 'Hardware Unboxed'.



**AMD**  
Ryzen 9  
3900X 12-  
core, 24-  
thread  
unlocked  
desktop  
processor  
with  
Wraith  
Prism LED  
Cooler

★★★★★ (8843)

**Brand:** AMD  
**CPU Model:** Ryzen 9 3900X  
**CPU Speed:** 4.6 GHz  
**CPU Socket:** Socket AM4  
**Processor Count:** 12



# Products



[OVERVIEW](#)

[SPECIFICATIONS](#)

[DRIVERS & SUPPORT](#)

## AMD Ryzen™ 9 3950X

### Specifications

**# of CPU Cores:** 16

**# of Threads:** 32

**Base Clock:** 3.5GHz

**Max Boost Clock** ⓘ: Up to 4.7GHz

**Total L1 Cache:** 1MB

**Total L2 Cache:** 8MB

**Total L3 Cache:** 64MB

**Unlocked** ⓘ: Yes

**CMOS:** TSMC 7nm FinFET

**Package:** AM4

**PCI Express® Version:** PCIe 4.0  
x16

**Thermal Solution (PIB):** Cooler  
Not Included, Liquid Cooling  
Recommended

**Default TDP / TDP:** 105W

### System Memory

**System Memory Specification:**  
Up to 3200MHz

**System Memory Type:** DDR4

**Memory Channels:** 2

### Foundation

**Product Family:** AMD Ryzen™  
Processors

**Product Line:** AMD Ryzen™ 9  
Desktop Processors

**Platform:** Boxed Processor

**OPN Tray:** 100-000000051

**OPN PIB:** 100-100000051WOF

# Products

## Ryzen vs Threadripper



**David Rho**, Partner at MMG Partners (2001-present)

Answered November 20, 2019



### What are the differences between the Ryzen 9 3950X and Threadripper?

The main differences between Ryzen 9 3950X and (I assume the new) Threadripper is cores, PCIe lanes, Cache, and Memory bandwidth.

The Ryzen 9 3950X has 24 PCIe lanes of which 4 are to the chipset. Now I would like to point out that these are PCIe 4.0 lanes which are twice as fast as PCIe 3.0 so this is going to help get data in and out of the CPU much faster. The 3950X also has 16 cores and 32 threads spread out over 2 chiplets with a third chiplet acting as the IO die. And the 3950X supports dual channel memory configurations. This can become a constraint on the CPU as you need to feed so many cores but have limited memory bandwidth. The 3950X has 64MB of L3 cache.

The new Threadripper 3960X starts with 24 cores and the 3970X has 32 cores. These are spread out over 5 chiplets with 4 being the CPU cores and the 5th being the IO die. The CPU supports a total of 72 usable PCIe 4.0 lanes. These CPUs support quad channel memory and the rumor is that there will be another Threadripper CPU and Chipset the supports 8 channel memory. The 3960X has 140MB of cache and the 3970X has 144MB of cache, but that includes both L2 and L3 cache but doing the math, there is 128MB of L3 cache for each CPU.


The TL;DR is that Threadripper is more scalable than Ryzen, but then again the Threadripper is more expensive (CPU and MB) and the chips target slightly different markets. Think of the 2950X as a baby Threadripper that also works well in gaming (based on the reviews).

# Products

## Ryzen 9

### How many transistors are in a Ryzen 9?

You have my answer on Intel i9 here: [How many transistors are in core i9? - Quora](#)

Googling I found this: [Ryzen - Wikipedia](#) 

AMD is a bit of "cheating" cause they use several dies instead of one big like Intel. AMD Ryzen is very simple. It always has IO die having 2.09 billion transistors and then number of CCDs (chiplet) each with two CCX and each CCX has up to 4 cores - 8 cores per CCD or chiplet.

Zen2 CCDs have up to 3.8 billion while Zen3 has up to 4.15 billion transistors.

I say up to cause some cores are disabled (malfunctioning).

So, Ryzen 9 3950x has 9.69 billion transistors while 5950x has 10.39 billion.

Ryzen 9 is nice chip to show how binning is important. Binning is process in fab when wafer is tested and various bins created. AMD manufacturers one or few identical wafers having CCDs but after testing CCDs do not behave identically, some support higher frequencies then other. Also some consume less power then others. Notice this, 3900 has 12 cores, 3.8/4.6 GHz and is 105W. At same time 3950 has 16 cores, 3.5/4.7 GHz and is also 105W. That's binning!

When someone buys CPU he/she plays silicon lottery. It might get CPU which operates at specified frequencies but could also get one much better working on higher frequencies.

Identical is for memories, DDR. There are no 4600, 5000, 6000 or even 7000 DDR. They are "same" dies but memory manufacturers test each one, bin them, and use best ones for crazy DDRs.

# Products

## AVX



**Joe Zbiciak**, I have been programming since grade school

Answered Sat



The **VEX** [↗](#) prefixes are used to encode the AVX and AVX2 vector instruction set extensions.

The **EVEX** [↗](#) prefix is used to encode the AVX512 vector instruction set extensions.

You don't typically ask for a VEX or EVEX prefix explicitly. Rather, the assembler emits prefixes as necessary when encoding instructions that require them.

Compilers will generate vector instructions when autovectorizing, or when the code explicitly requests vector instructions. You likely need to enable a certain level of optimization before the compiler will attempt autovectorizing. You likely also need to tell it to enable support for particular sets of instruction set extensions. Details vary by compiler.

Vector instructions speed up computation by performing many identical computations in parallel. These usually show up in data processing loops.

I realize that's vague. It's purposefully so. I've used vector techniques for graphics, matrix arithmetic, string processing, signal processing, cryptography, and so on. What all these applications have in common is that they loop over large quantities of data performing identical operations to many consecutive elements.

## Current product lines [\[ edit \]](#)

### CPU and APU products [\[ edit \]](#)

AMD's portfolio of CPUs and APUs as of 2020

- **Athlon** - brand of entry level CPUs (Excavator) and APUs (Ryzen)
- **A-series** - *Excavator* class consumer desktop and laptop APUs
- **G-series** - *Excavator* and *Jaguar* class low power embedded APUs
- **Ryzen** - brand of consumer CPUs and APUs
- **Ryzen Threadripper** - brand of prosumer/professional CPUs
- **R-series** - *Excavator* class high performance embedded APUs
- **Epyc** - brand of server CPUs
- **Opteron** - brand of microserver APUs<sup>[229]</sup>

### Graphics products [\[ edit \]](#)

AMD's portfolio of dedicated [graphics processors](#) as of 2017

- **Radeon** – brand for consumer line of graphics cards; the brand name originated with ATI.
  - **Mobility Radeon** offers power-optimized versions of Radeon graphics chips for use in laptops.
- **Radeon Pro** – Workstation Graphics card brand. Successor to the [FirePro](#) brand.
- **Radeon Instinct** - brand of server and workstation targeted machine learning and [GPGPU](#) products

### Radeon-branded products [\[ edit \]](#)

#### RAM [\[ edit \]](#)

In 2011, AMD began selling Radeon branded [DDR3 SDRAM](#) to support the higher bandwidth needs of AMD's APUs.<sup>[230]</sup> While the RAM is sold by AMD, it was manufactured by [Patriot Memory](#) and VisionTek. This was later followed by higher speeds of gaming oriented DDR3 memory in 2013.<sup>[231]</sup> Radeon branded [DDR4 SDRAM](#) memory was released in 2015, despite no AMD CPUs or APUs supporting DDR4 at the time.<sup>[232]</sup> AMD noted in 2017 that these products are "mostly distributed in Eastern Europe" and that it continues to be active in the business.<sup>[233]</sup>

## Technologies [\[ edit \]](#)

---

### CPU technologies [\[ edit \]](#)

As of 2017 technologies found in AMD CPU/APU products include:

- **HyperTransport** - a high-bandwidth, low-latency system bus used in AMD's CPU and APU products
- **Infinity Fabric** - a derivative of HyperTransport used as the communication bus in AMD's [Zen microarchitecture](#)

### Graphics technologies [\[ edit \]](#)

As of 2017 technologies found in AMD GPU products include:

- **AMD Eyefinity** – facilitates [multi-monitor](#) setup of up to 6 monitors per graphics card
- **AMD FreeSync** – display synchronization based on the [VESA Adaptive Sync](#) standard
- **AMD TrueAudio** – acceleration of audio calculations
- **AMD XConnect** – allows the use of [External GPU](#) enclosures through [Thunderbolt 3](#)
- **AMD CrossFire** - multi-GPU technology allowing the simultaneous use of multiple GPUs
- **Unified Video Decoder (UVD)** – acceleration of video decompression (decoding)
- **Video Coding Engine (VCE)** – acceleration of video compression (encoding)

### Software [\[ edit \]](#)

- **AMD Catalyst** is a collection of [proprietary device driver](#) software available for [Microsoft Windows](#) and [Linux](#).
- **AMDGPU** is AMD's [open source device driver](#) supporting the [GCN](#) architecture, available for Linux.



- A partnership between AMD and Alpha Processor Inc. developed [HyperTransport](#), a point-to-point interconnect standard which was turned over to an industry standards body for finalization.<sup>[254]</sup> It is now used in modern motherboards that are compatible with AMD processors.
- AMD also formed a strategic partnership with IBM, under which AMD gained [silicon on insulator](#) (SOI) manufacturing technology, and detailed advice on [90 nm](#) implementation. AMD announced that the partnership would extend to 2011 for [32 nm](#) and [22 nm](#) fabrication-related technologies.<sup>[255]</sup>
- To facilitate processor distribution and sales, AMD is loosely partnered with end-user companies, such as [HP](#), [Dell](#), [Asus](#), [Acer](#), and [Microsoft](#).<sup>[256]</sup>
- In 1993, AMD established a 50–50 partnership with [Fujitsu](#) called FASL, and merged into a new company called FASL LLC in 2003. The joint venture went public under the name [Spansion](#) and ticker symbol SPSN in December 2005, with AMD shares drop to 37%. AMD no longer directly participates in the Flash memory devices market now as AMD entered into a non-competition agreement, as of December 21, 2005, with Fujitsu and Spansion, pursuant to which it agreed not to directly or indirectly engage in a business that manufactures or supplies standalone semiconductor devices (including single chip, multiple chip or system devices) containing only Flash memory.

# Section

---



# Ryzen CPU's on Zen 3

# AMD CES 2023



Jan 2023

**AMD @ CES 2023**

LAS VEGAS, NEVADA  
JANUARY 4TH, 6:30 PM PT  
[AMD.COM/CES](https://www.amd.com/ces)

**AMD**  
together we advance

The image is a promotional poster for the AMD event at CES 2023. It features a dark background with a glowing, stylized AMD logo on the left. The text is centered and right-aligned, providing event details. The AMD logo and tagline are in the bottom right corner.

# AMD CES 2023

Jan 2023


## AMD's new Ryzen 7000 mobile processors include a massive 16-core chip

Story by Monica Chin • Yesterday 7:30 PM



Comments


This may seem like an underdog compared to [Intel's top 13th-Gen chips](#), which have 24 cores — Intel has claimed that its Core i9-13980HX is the “world’s fastest mobile processor”. This chip, however, only has eight performance cores and 16 efficiency cores, while all 16 of the Ryzen’s will be going full-speed ahead. This is essentially two full eight-core chips stuffed into one.



# AMD CES 2023

Jan 2023

## Announcing AMD RYZEN™ 7045HX SERIES



- Processor Architecture
- 5nm Technology
- RDNA™ 2 Graphics
- up to 16 Cores 32 Threads
- up to 5.4GHz Boost
- up to 80 MB L2+L3 Cache
- 55W+ TDP


Zen4

5nm

16 P


## AMD RYZEN™ 7040 SERIES

Core i7 1280P | Apple | Ryzen™ 9 7940HS



# AMD CES 2023

Jan 2023



Announcing  
**AMD RYZEN™ 9 7950X3D**  
Ultimate processor for gamers and creators

- ZEN 4 Processor Architecture
- 5nm Technology
- up to 16 Cores 32 Threads
- up to 5.7GHz Boost
- up to 144 MB L2+L3 Cache
- 120W TDP

Zen4

5nm

16 P

RDNA3



Announcing  
**AMD RADEON™ RX 7600M XT**

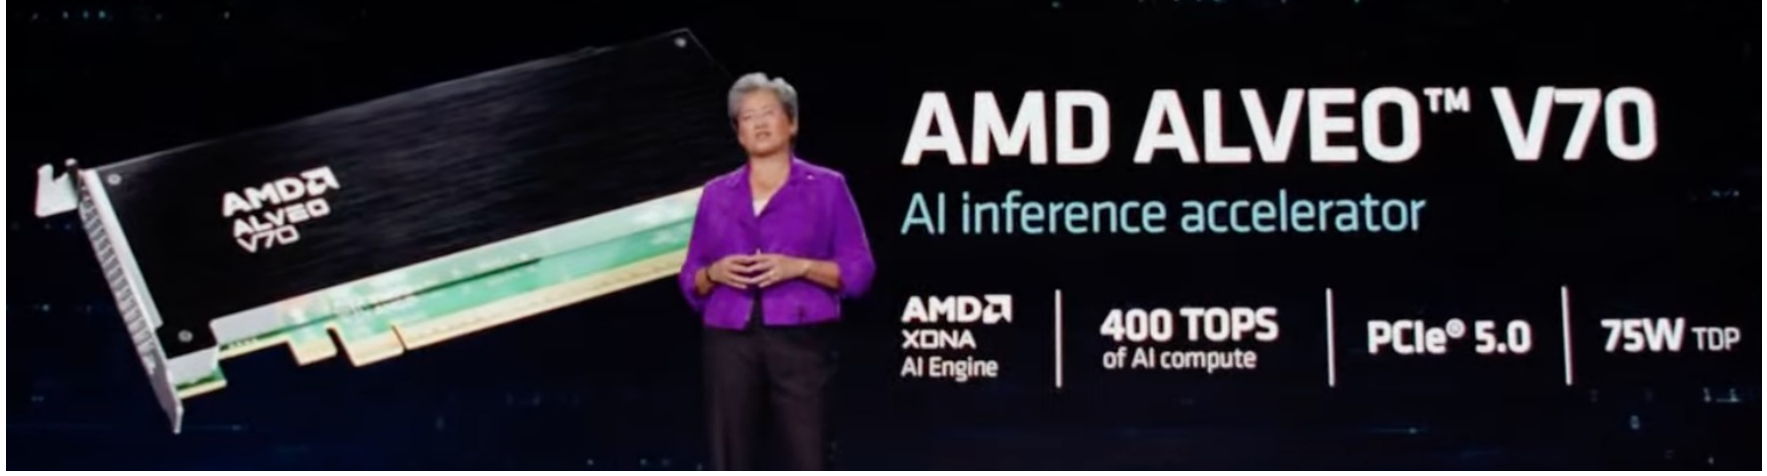
- 32 RDNA™ 3 Compute Units
- 8 GB GDDR6
- 128-Bit Memory Bus
- 75-120W TDP
- 6nm Process Technology

32 CU

6nm

# AMD CES 2023

Jan 2023



**AMD ALVEO™ V70**  
AI inference accelerator

**AMD**  
XDNA  
AI Engine

**400 TOPS**  
of AI compute

**PCIe® 5.0**

**75W TDP**

A woman in a purple jacket stands next to a large AMD Alveo V70 AI inference accelerator card. The card is black with green and gold accents and has the AMD Alveo V70 logo on it.



**AMD** **NASA**

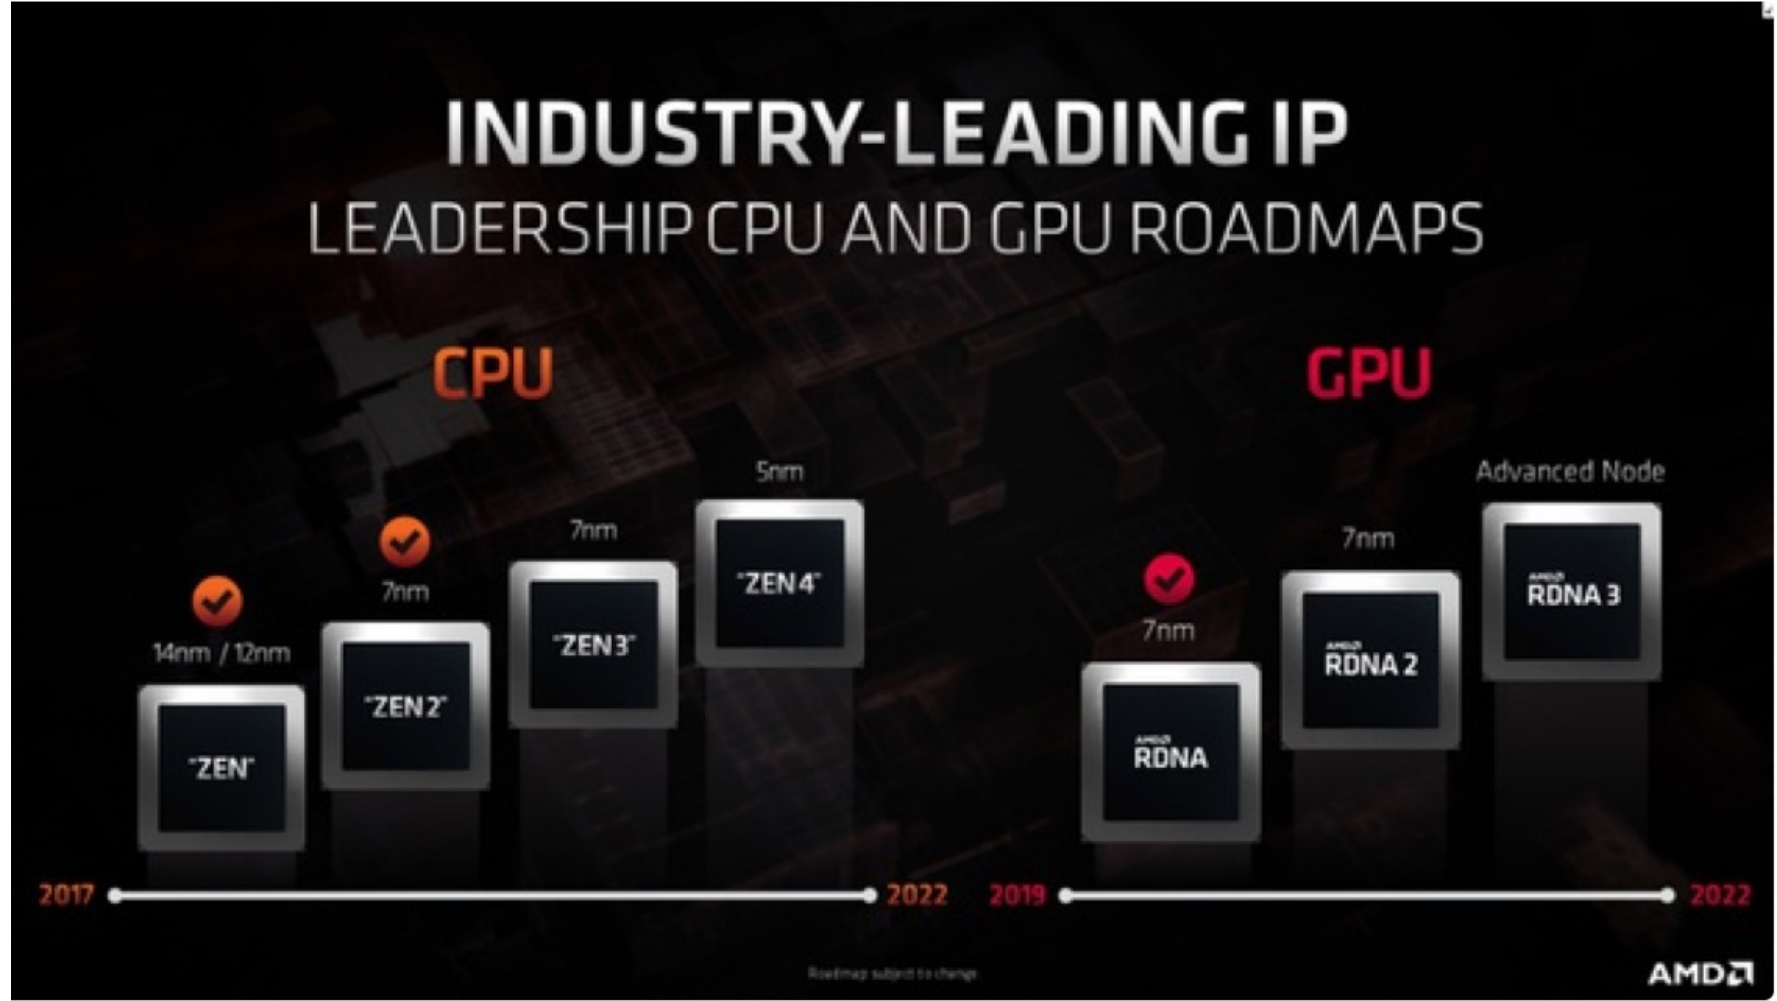
**AMD Xilinx FPGAs**  
powering Artemis and Orion

**ARTEMIS** **ORION**

Two women, one in a blue NASA jacket and one in a purple jacket, stand on a stage in front of a large screen displaying the Earth from space. The screen features the AMD and NASA logos, and the text 'AMD Xilinx FPGAs powering Artemis and Orion'. Below the text are the Artemis and Orion logos.

# AMD Process Roadmap

When will AMD be transitioning to 5nm like Apple M1?





# AMD Process Roadmap

## HIGH PERFORMANCE MOMENTUM

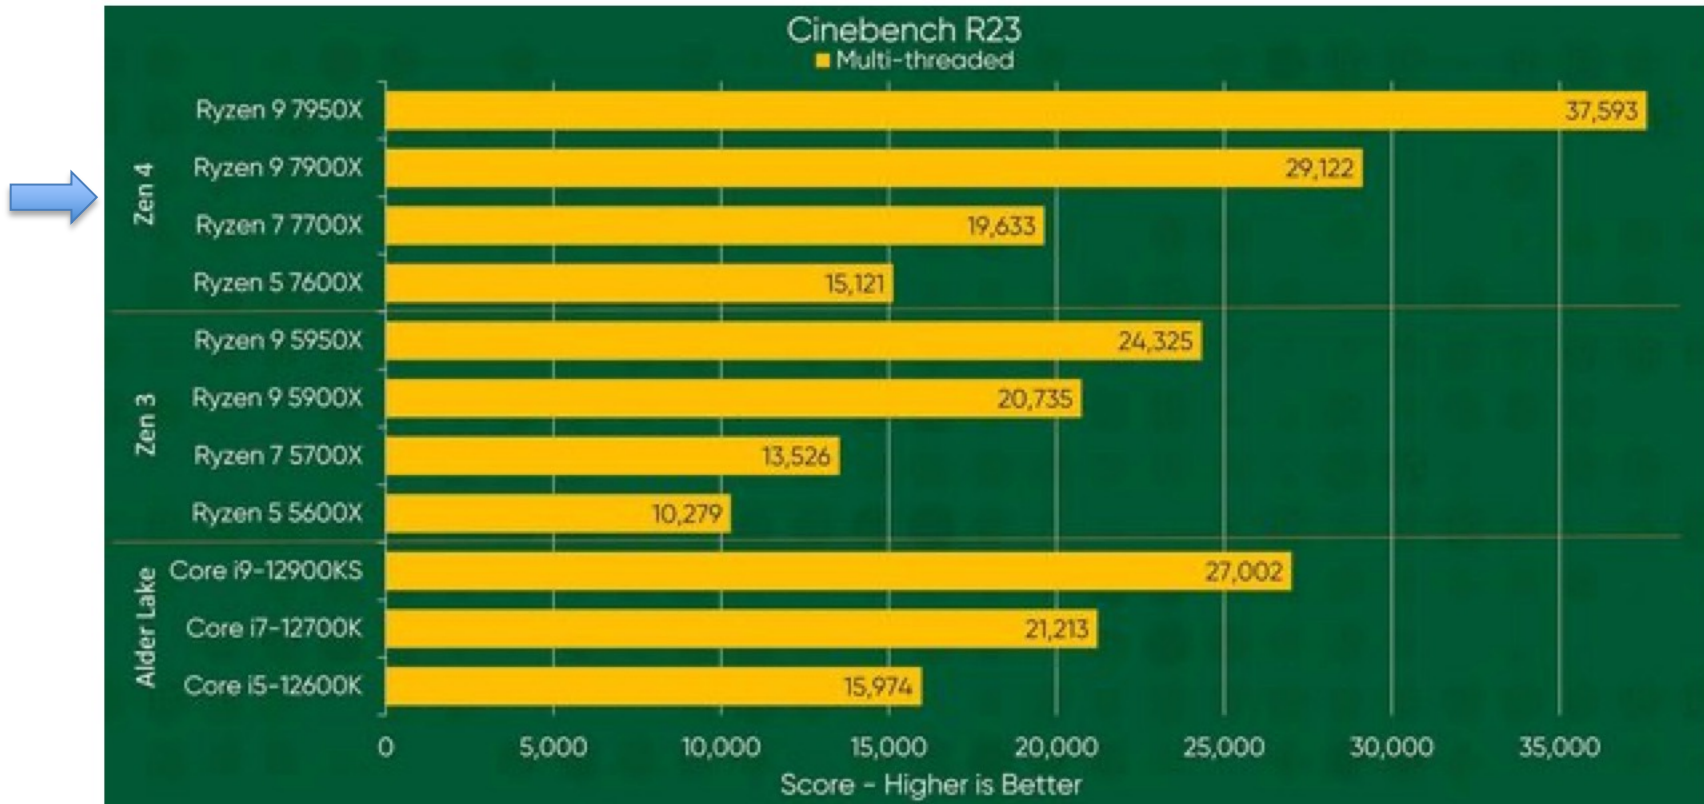


2017

2022


# Zen 4 Benchmark

- Multi-threaded - here Ryzen leads cause it has more cores



➤ **Zen 4** is 50% better than Zen 3

# AMD Zen 3 CCX



- Increased 4→8 cores per CCX
- Same max. 4MB of L3 per core
- Single CCX per chiplet
- 2x L3 cache directly accessible per core

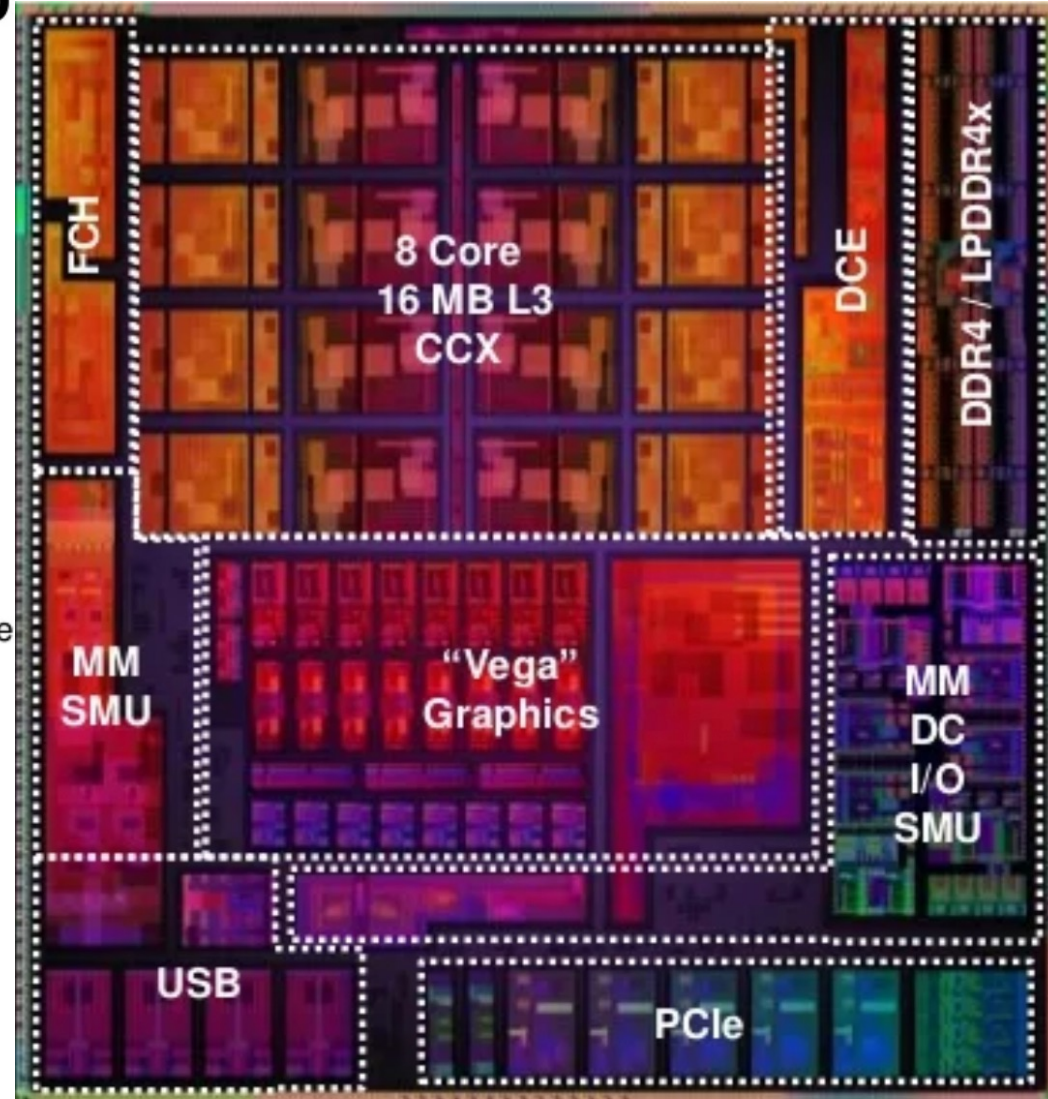
**“Zen 3”: AMD 2nd Generation 7nm x86-64  
Microprocessor Core** 2022

COMP122

## APU Monolithic Chip

### Key chip features


- 8 core, 16MB L3 CCX
- 8 compute-unit (“Vega”) graphics
- 2 memory controllers
  - DDR4 up to 3200 MT/s
  - LPDDR4x up to 4266 MT/s
- Multimedia (MM) engines
  - 2<sup>nd</sup> Gen Video Codec<sup>1</sup>, 3<sup>rd</sup> Gen Audio ACP
- 2<sup>nd</sup> Gen display controller (DC)
- I/O controllers
  - PCIe® Gen4, USB-C, USB-3.1, USB-2.0, NVMe
- System management unit (SMU)
- Fusion controller hub (FCH)



# AMD 3D V-Cache

## AMD 3D V-Cache Ready

- Two columns of TSVs on left/right side of the L3 cache
- AMD 3D V-Cache extends L3 Cache capacity by 64MB (3x)
- Total inter-die bandwidth: >2 TB/s
- All control and routing to the cores is implemented on the base die, so AMD 3D V-Cache can be completely focused on density



**“Zen 3”: AMD 2nd Generation 7nm x86-64  
Microprocessor Core**

<sup>1</sup>AMD, Fort Collins, CO, <sup>2</sup>AMD, Santa Clara, CA, <sup>3</sup>AMD, Austin, TX

# AMD64

## **AMD64 Architecture Programmer's Manual: Volumes 1-5**

### **Overview of the AMD64 Architecture**

---

The AMD64 architecture is a simple yet powerful 64-bit, backward-compatible extension of the industry-standard (legacy) x86 architecture. It adds 64-bit addressing and expands register resources to support higher performance for recompiled 64-bit programs, while supporting legacy 16-bit and 32-bit applications and operating systems without modification or recompilation. It is the architectural basis on which new processors can provide seamless, high-performance support for both the vast body of existing software and 64-bit software required for higher-performance applications.

The need for a 64-bit x86 architecture is driven by applications that address large amounts of virtual and physical memory, such as high-performance servers, database management systems, and CAD tools. These applications benefit from both 64-bit addresses and an increased number of registers. The small number of registers available in the legacy x86 architecture limits performance in computation-intensive applications. Increasing the number of registers provides a performance boost to many such applications.

# Zen uArch

## AMD THE EVOLUTION OF "ZEN"

RESOURCE	"ZEN"	"ZEN 2"	"ZEN 3"
Issue width	10	11	16
INT reg	168	180	192
INT sched	84	92	96
FP reg	160	160	160
ROB	192	224	256
FADD, FMUL, FMA	3/4/5	3/3/5	3/3/4
FP width	128	256	256
L1 BTB	256	512	1024
L2 BTB	4k	7k	6.5k

### CORE COMPARISON

# Zen uArch



## THE EVOLUTION OF "ZEN"

RESOURCE	"ZEN"	"ZEN 2"	"ZEN 3"
LDQ	72	72	72
STQ	44	48	64
Micro-Op-cache	2k	4k	4k
L1 Icache	64k	32k	32k
L1 Dcache	32k	32k	32k
L2 cache	512k	512k	512k
L3 cache/core	2M	4M	4M
L2 TLB size	1.5k	2k	2k
L2 TLB latency	8	6	6
L2 latency	12	12	12
L3 latency	35	39	46

### CACHE COMPARISON



# New AMD Ryzen 5

## COMP122 UserBenchmark

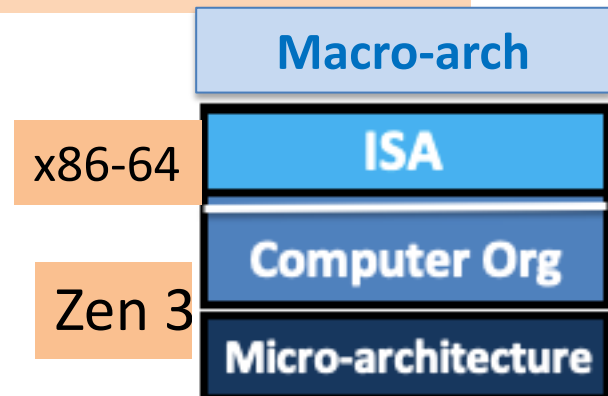
### AMD Ryzen 5 3600 \$199

AMD's Ryzen 5 3600 is a 6-core, 12-threaded processor which succeeds the Ryzen 5 2600 **improving upon it by 13% in terms of overclocked performance**. The 3600 is in competition with Intel's 6-core i5-9600K. AMD continues to push the multi-core performance envelope: benchmarks show that the 3600 has a 27% overclocked 64-core lead over the 9600K but that the **i5-9600K leads by 14% on single to hex core workloads which translates to 10% higher EFps in most of the today's top games (e.g. PUBG, GTAV and CSGO)**. Additionally, the 3600's memory controller, although significantly improved over previous Ryzen iterations, still has limited bandwidth and high latency which adversely impacts gaming. Weaknesses in memory architecture are not readily picked up by CPU benchmarks but they are apparent whilst gaming. Cheaper CPUs such as the **9400F deliver better gaming performance in nearly all of today's popular games**. At \$190 USD, the

➤ vs. Intel Core i5-9600K



6 CPU cores (SMT-12)



# New AMD Ryzen 5

Zen 3

6 CPU cores (SMT-12)



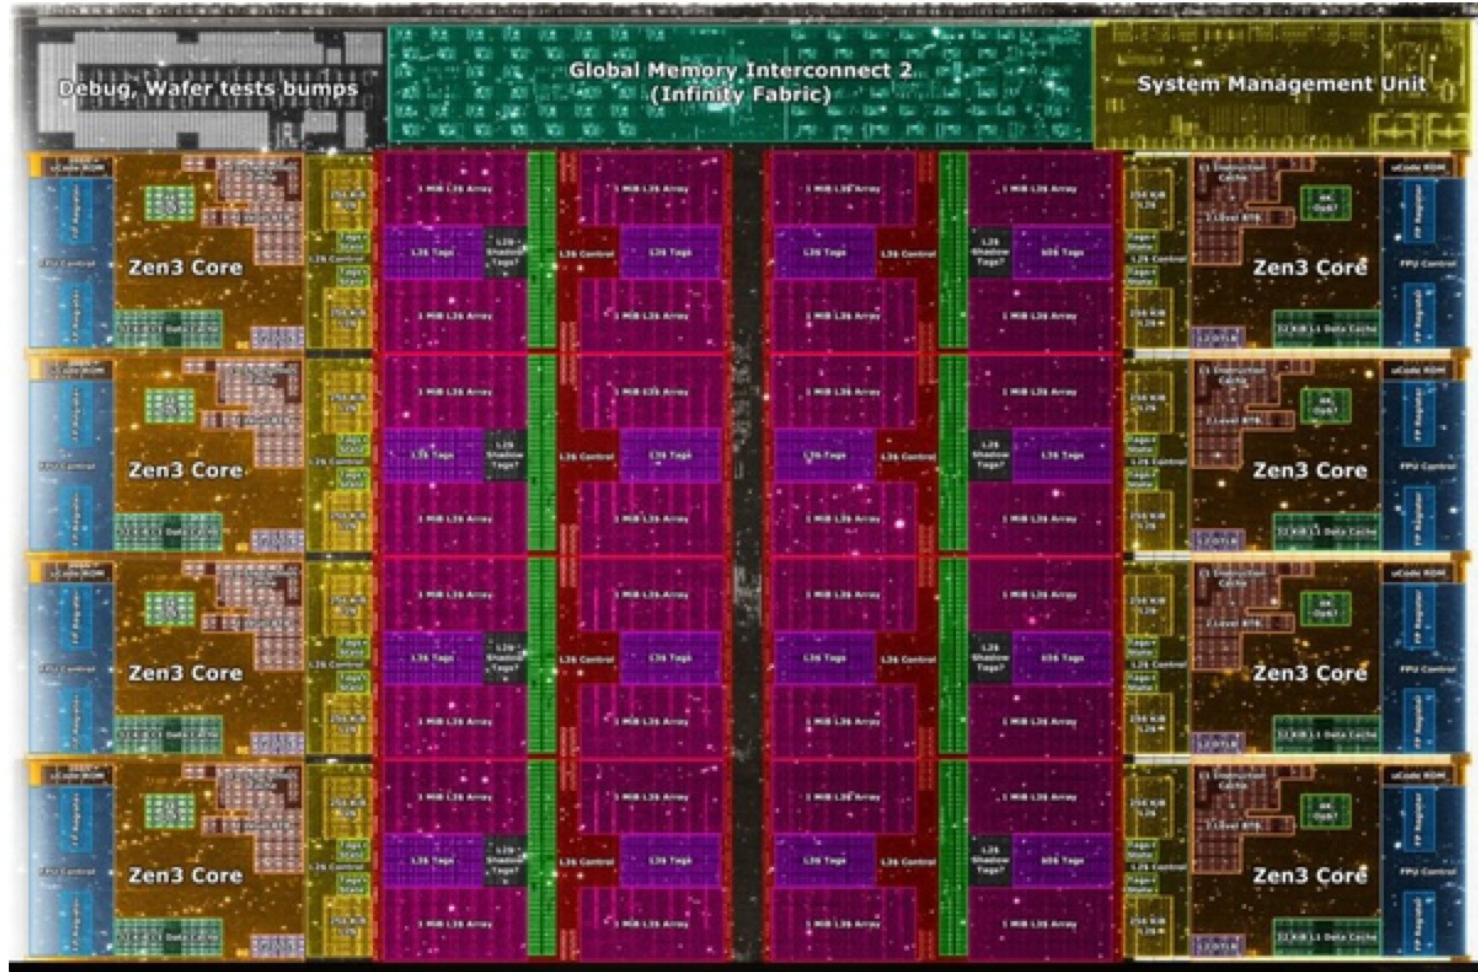
## AMD Ryzen 5 5600X \$300 🔥

The Ryzen 5 5600X is both the entry-level and best value for money 5000 series CPU. The 5600X is a hex-core 12 thread processor with a base clock speed of 3.7 GHz boosting to 4.6 GHz. It has 35 MB of cache and a TDP rating of 65W. A cooler is included in the MSRP of \$300 USD, but cheap after-market coolers (such as the [\\$20 GAMMAXX 400](#)) are *far more effective* and therefore worth the upgrade. Notably, AMD's new Zen 3 architecture has vastly improved single-core performance and lower memory latency, which leads to a *significant effective speed advantage over its predecessor, the 3600X*. Last year, AMD's marketers secured significant sales of the 3000 series CPUs despite a 15% performance deficit against lower priced Intel parts. The games, specific scenes,

# New AMD Ryzen 5

Zen 3

This is Zen 3 (Ryzen 5000) die:



Core, without cache, is small part with text Zen3 Core on it. Everything else are various caches.

*AMD News*  
**AMD's new Zen Processor**

AMD made radical alterations to its Zen design while keeping itself distant from an ugly past. The company knew it had to make the changes to become a force to reckon with in the server and PC markets. So when the designers of the chip sat down to map the Zen design, they had two priorities: To boost CPU performance to maximum and to stabilize power efficiency.

According to a company spokesperson, the chips will come with 8 to 32 cores. The 32-core chips may come in the quad-CPU configurations although those details haven't been finalized yet.



AMD Zen  
*stock news usa (2016)*

(Click image to view full size)


**5-19B** Transistors

- ❖ CPU performance
- ❖ Power efficiency

**8-32** cores


Source: *stocknewsusa (2016-08-26)*

[Inside AMD's Production Of The Zen CPU](#)



Ryzen 7 will have three CPUs to start, all having eight cores and supporting simultaneous multi-threading:

- Ryzen 7 1800X: 8C/16T, 3.6 GHz base, 4.0 GHz turbo, 95W, \$499
- Ryzen 7 1700X: 8C/16T, 3.4 GHz base, 3.8 GHz turbo, 95W, \$399
- Ryzen 7 1700: 8C/16T, 3.0 GHz base, 3.7 GHz turbo, \$329



Ryzen

4.8 billion transistors and more than 2,000m of signal wire



# AMD Ryzen



**Norman Latifov**, knows Turkish

Answered 2h ago



It is the latest mobile CPU from AMD. Ryzen 4000 CPUs are only available on laptops and they are the fastest mobile CPUs available right now. I am using a Lenovo Yoga slim 7 with r7 4800u and before I bought it I did a lot of research. Based on reviews, they are even faster than 11th gen Intel CPUs that are yet to come. Although their integrated GPU is not as good as iris graphics (Intel 11th gen CPUs' integrated GPU), vega series GPUs are still a good option. And in my opinion, ryzen 4000 CPUs have a great multicore performance.



**Zachary Hawkshaw**, AMD Hardware Connoisseur

Answered 10h ago



I don't know where you got your information, but that's not true. The 2700X only has 4.8 billion transistors while the 3700X has 19.2 billion.

Transistor count doesn't necessarily mean more performance by itself. The 3700X is better because it has a newer architecture (Zen2 vs Zen+) with improvements to the Infinity Fabric, as well as a higher turbo frequency.

# AMD Ryzen



**John B. Anderson III**, IT Consultant, PC Integrator, 20+ Years in IT and Gaming.



Answered Mon

Well, it's only on the Laptop side... Since they skipped it to make them both match desktop/laptop for Zen 3 architecture.

Laptops with the 4XXX name are actually Zen 2 processors.

See some guy in marketing thought it'd be a good idea to call Zen processors on laptops 2XXX series, and so when the 2XXX series came out on desktop the laptop was already at 2XXX so they called them 3XXX.

Example Desktop CPU Ryzen 5 2600x the laptop version would be a Ryzen 5 3550h

[AMD Ryzen 5 2600 vs 3550H](#) with the typical lower performance on laptop vs desktop CPU. Same Zen+ architecture in both.

On the Zen 2 Architecture

[AMD Ryzen 5 3600 vs 4600HS](#) we see the Ryzen 5 4600 HS ~5% of the speed.

There are no Zen 3 laptop processors as of the time I'm writing 11/2020

Typically we call it a "Gen / Generation" when they name a model with a change in the first digit.

Example Core i5 10XXX would be a 10th gen i5. Ryzen 7 1800x would be a 1st Gen Ryzen, However, Since the laptops, 4XXX were the only "Gen 4s" but they were technically Gen 3's well the answer is somewhat tricky.



# AMD

CPU

128GiB

DRAM

## Consumer Processors Processors for Desktops

### AMD Ryzen™ Threadripper Processors

For professionals, serious content creators, and elite enthusiasts who need the most powerful desktop processors in the world!

- From 8 to 64 cores
- From 16 to 128 processing threads



The big old Threadripper in my desktop PC only supports 128GiB DRAM. But processors built for servers, Intel® Xeon® and AMD EPYC, will support 1TB+ DRAM. The EPYC 7742 will support 4TiB DRAM, and the price of the DRAM won't seem so shocking when you learn the price of your 64-core CPU is about \$7,500!






# AMD



AI CPU June 2024







## AMD Ryzen™ Processors and AMD Ryzen™ Processors with Radeon™ Graphics

For desktop enthusiasts, demanding gamers, and content creators who need a high-performance PC

- From 4 to 16 cores
- Up to 32 processing threads
- Some models include Radeon™ graphics



## AMD Athlon™ Processors with Radeon™ Vega Graphics

For entry-level users who value advanced technology, fast responsiveness, and the power to handle graphics card upgrades. Now including the new unlocked Athlon™ 3000G.<sup>2</sup>

- 4 processing threads
- Includes Radeon™ Vega graphics
- Advanced 'Zen' processor technology



# AMD

GPU



### AMD Radeon™ Boost

Turbocharge your game

[LEARN MORE](#)

### AMD Radeon™ Anti-Lag

Delivering even faster click-to-response times

## Technologies for Gaming

### AMD Enhanced Sync Technology

Now available for games based on DirectX® 9, 11, 12 and Vulkan® APIs as well as on all GCN-based GPUs, GCN-based GPU combinations, and/or AMD Eyefinity Technology display combinations.

[LEARN MORE](#)



### AMD Virtual Super Resolution



### AMD TressFX Hair

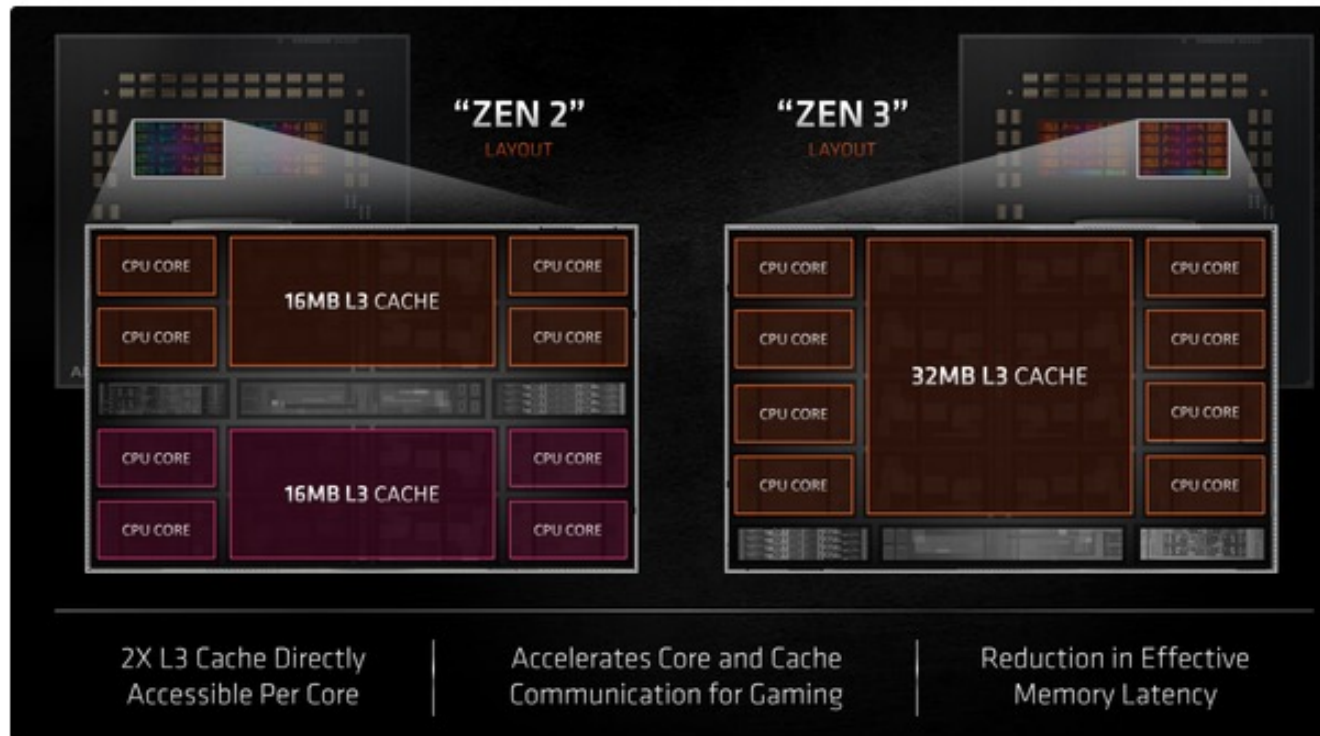


### The Vulkan® API

# AMD Zen L3 Cache

Quora Home Following Answer 186 Spaces Notifications 763 Search

One area that AMD has lagged behind Intel over the lifetime of the Zen brand is in gaming performance. It's no secret that in the company's push to lower the cost per core of its flagship processors (through the introduction of chiplet-based architectures), the design decisions have resulted in more latency between core complexes. That manifests itself in reduced performance in certain PC gaming scenarios--especially at the favored 1080p resolution used by most gamers.



This is down to how chips are designed, and, more specifically, how they're laid out on

Another way that AMD looks poised to continue its winning streak over Intel is in its platform and required socket adoption for new PCs, PC builds, and upgrades...or rather, its of required adoption. Instead of forcing buyers onto a new motherboard platform with a new style of CPU socket every other generation of chips (the typical cadence in recent years for Intel's desktop processors), Zen 3 will mark the third launch in the Zen line to feature some level of compatibility with motherboards based on the now-venerable Socket AM4.

### BIOS UPDATES FOR RYZEN™ 5000 SERIES

#### GETTING YOUR MOTHERBOARD DROP-IN READY



#### AMD 500 Series Chipsets

- **GET READY:** AMD 500 Series motherboards require a BIOS with AGESA 1.0.8.0 (or newer) for POST/boot (already available)
- **AT LAUNCH:** Users should upgrade to a BIOS with AGESA 1.1.0.0 (or newer) for the best experience on November 5

#### AMD 400 Series Chipsets

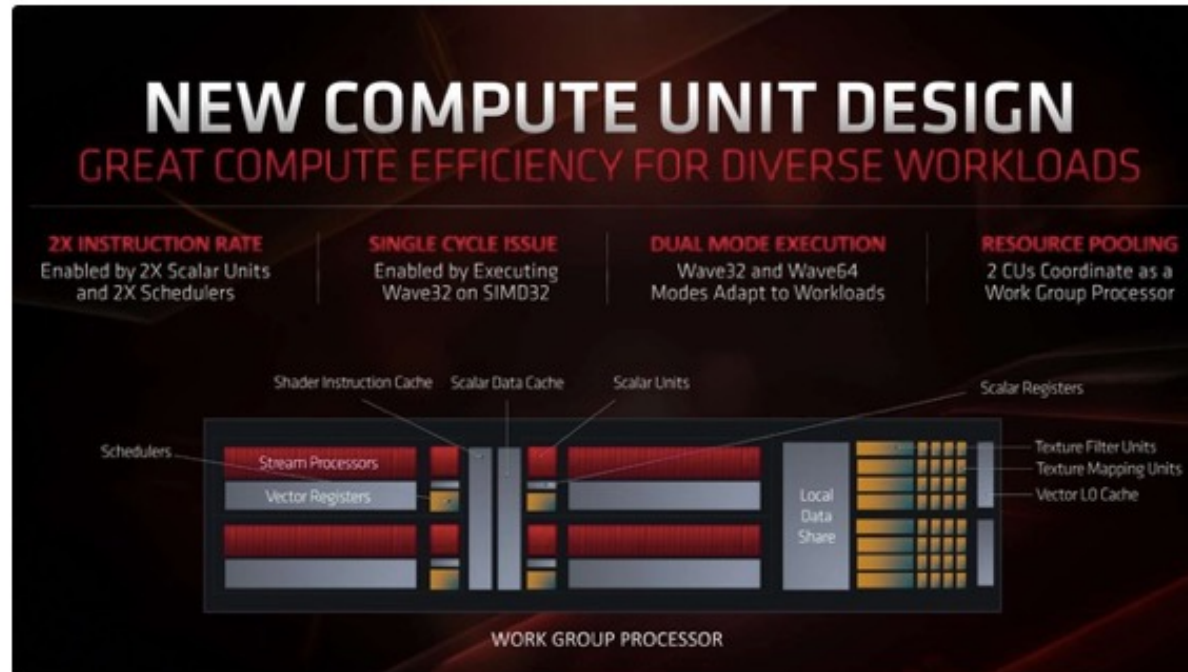
- BIOS updates for AMD Ryzen 5000 Series processors currently in development with motherboard partners
- Customers should expect first beta releases for AMD 400 Series motherboards starting in January, 2021

The one caveat? Unlike Zen 2, which is compatible with just about AM4-based motherboard, the cutoff for Zen 3 is a bit higher up the chipset stack this time. The new CPUs will work only with motherboards from the X470, B450, and later chipset generations. (That includes the new X570 and B550 boards.) Plus, it's down to motherboard manufacturers to make it work, issuing the proper updates.

Back in May, the company clarified Zen 3 AM4 compatibility, claiming that a BIOS update would be required for any users of either X470 or B450 motherboards. Now, we

# AMD Zen Cores


Each core (DCU) has four 32-lane SIMD units, so that it each core can perform 128 32-bit floating point multiply-accumulations(FMAs) per clock cycle and the whole chip 5120 32-bit FMAs per clock cycle.



But each Zen2 or Zen3 CPU core also has 2 256-bit (8-lane for 32-bit) SIMD FMA units, so that each core is capable of calculating 16 32-bit FMAs per clock cycle, so for 8-core CPU such as 3800X or 5800X, the whole chip can perform 128 32-bit FMAs per clock cycle.

So, the core count difference between CPUs and GPUs is not many hundreds of times in reality, it's less than 10 times, and the difference in parallel execution units is only about 20-40 times, not hundreds of times.


# AMD Zen Cores



Each core (DCU) has four 32-lane SIMD units, so that it each core can perform 128 32-bit floating point multiply-accumulations(FMAs) per clock cycle and the whole chip 5120 32-bit FMAs per clock cycle.

# AMD Zen Cores


To make the CPUs execute the serial code as quickly as possible, they have very big and complex cores which have things like very advanced branch predictors and huge caches to minimize time wasted for stalls. The actual execution units are only very small part of the transistor count or area of the core.



Here is a picture of AMD Zen core. The very small part "ALU" consists all the execution units for most commonly used integer instructions, and FPU is mostly the execution units of the floating point and SIMD instructions. L1D\$, L1I\$ and L2\$ are cache memory, BPU is branch prediction unit which just tries to guess what the core should do next to minimize stalls.

# New HPC Products

## Epyc 7000







# Epyc

March 2021



# Epyc

March 2021

## OUR FOCUS

### HIGH-PERFORMANCE COMPUTING SOLUTIONS



AMD EPYC™



AMD INSTINCT™




AMD RYZEN™  
DESKTOP



AMD RYZEN™  
MOBILE



CONSOLE  
GAMING



AMD RADEON™



# Epyc 7003



March 2021

## AMD EPYC™ 7003 FAMILY

AVAILABLE TODAY

CORES	MODEL
64 CORES	7763 7713/P
56 CORES	7663
48 CORES	7643
32 CORES	75F3 7543/P 7513
28 CORES	7453
24 CORES	74F3 7443/P 7413
16 CORES	73F3 7343 7313/P
8 CORES	72F3




### ALL-IN FEATURE SET

- 8-Channels of DDR4-3200
- 4TB Memory Capacity
- 128 Lanes PCIe® 4
- Infinity Guard Security Features
- Socket Compatible

# Epyc 7003

June 2024

## Introducing AMD EPYC™ 4004 Series Processors




# Epyc Chiplatets



## THE SECRET OF CHIPLATETS SUSTAINING AMD SOCKET AM4

<https://www.amd.com/en/technologies/zen-core-3>

- ▲ Chiplet transition enabled scalability up to 16 cores in AM4
- ▲ Allowed migration to "Zen 3" CCDs without disrupting the platform
- ▲ Enables in-place upgrades to "Zen 3" for AMD Socket AM4





March 2021

## News Story

### AMD Rolls Out Epyc 7003 Series Processors; Shares Rise

3:15 PM 3/15/2021 - MT Newswires

03:15 PM EDT, 03/15/2021 (MT Newswires) -- Advanced Micro Devices (AMD) said Monday that it has rolled out its new Epyc 7003 series CPUs. Shares of the Santa Clara, California-based company were up 1.8% in late-day trading. Price: 82.34, Change: +1.29, Percent Change: +1.59



## ENTERPRISE COMPUTING LEADERSHIP FROM DATA CENTER TO LAPTOP



### ENTERPRISE PCs

HP and Lenovo laptop and desktop enterprise portfolios powered by AMD Ryzen™ PRO Processors



### DATA CENTER PLATFORMS

AMD Instinct™ MI100  
The World's Fastest  
HPC GPU

Full portfolio of  
servers powered by  
AMD EPYC™ CPUs



### WORKSTATIONS

Enterprise workstations powered by  
AMD Threadripper™ PRO Processors  
and AMD Radeon™ PRO Graphics




# Epyc

March 2021



**~19% INSTRUCTIONS-PER-CLOCK INCREASE**  
COMPARED TO 2<sup>ND</sup> GEN AMD EPYC™

## “ZEN 3” CORE ARCHITECTURE



**8** CORE  
COMPLEX

### LOWER LATENCY

Significant Accelerator for HPC and Cloud Applications

**2X**

### DIRECT ACCESS L3 CACHE

Better Performance for Large Virtual Machines

**2X**

### THROUGHPUT FOR AI INFERENCE

Doubling of INT8 Pipeline and Faster Floating Point

**19%**

### CORE COMPLEX IPC UPLIFT

Performance on SPECrate® 2017 at Same Power



## FASTER IN CLOUD

BEST COMPUTE DENSITY FOR HYPERSCALE



3<sup>RD</sup> GEN AMD EPYC™

---

**106%**

FASTER THAN COMPETITION

---

2P INTEGER PERFORMANCE  
SPECrate@2017\_int\_base

## OUR “ZEN” JOURNEY

### “ZEN” / “ZEN+”

- › ~52% IPC
- › 4-core complex
- › 8MB L3 per complex
- › 14nm/12nm
- › Simultaneous multithread
- › SEV

### “ZEN 2”

- › ~15% IPC
- › 4-core complex
- › 16MB L3 per complex
- › 7nm
- › Chiplet design
- › FP-256
- › SEV-ES

### “ZEN 3”

- › ~19% IPC
- › New core layout
- › New cache topology
- › 7nm
- › Doubled INT8 throughput
- › New security features:  
Shadow Stack & SEV-SNP


2017

2020



March 2021


## LEADERSHIP ROADMAP



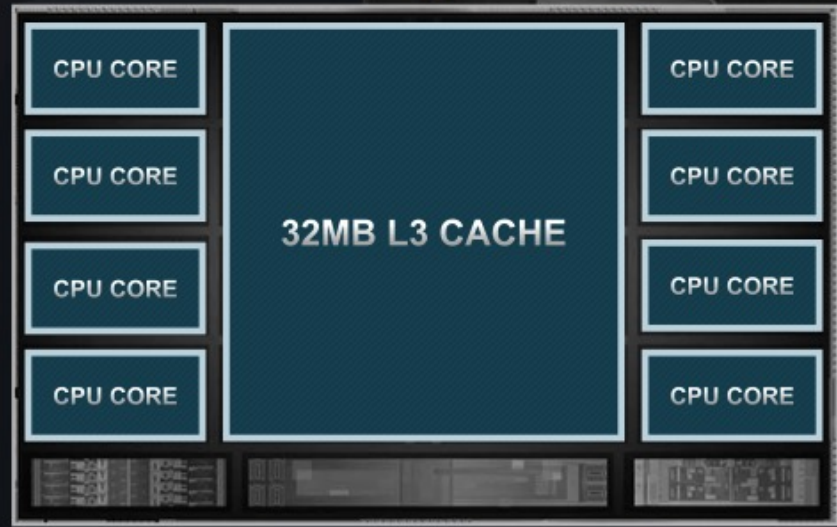
# Epyc

March 2021

**“ZEN 2”  
LAYOUT**



**“ZEN 3”  
LAYOUT**



2X L3 Cache Directly  
Accessible Per Core

Accelerates Core and Cache Communication  
for Enterprise and Cloud Workloads

Reduction in Effective  
Memory Latency




March 2021

## INDUSTRY LEADERSHIP

### "ZEN 3" ~19% IPC UPLIFT

**GEOMEAN OF 28 WORKLOADS**  
(FIXED 3.7GHZ FREQUENCY, 8 CORES)

**"ZEN 3" PERFORMANCE CONTRIBUTORS**



"ZEN 2"

"ZEN 3"

## FASTER IN HPC

ACCELERATING SCIENTIFIC BREAKTHROUGHS

**FP**



3<sup>RD</sup> GEN AMD EPYC™

---

**106%**

FASTER THAN COMPETITION

---

2P FLOATING POINT PERFORMANCE  
SPECrate@2017\_fp\_base

# AMD EPYC

# AMD

Servers

CPU


64 CORES    48 CORES    32 CORES    24 CORES    16 CORES    12 CORES

8 CORES

up to  
**128** threads  
base  
**2.6GHz**

up to  
**3.4GHz** boost  
cache  
**256MB**


X86 compatibility



2P Intel vs 1P EPYC comparison<sup>7</sup>




Model	2x6262V	1x7702P
Cores	48	64
Memory Capacity	2TB	4TB
Max Memory Frequency	2400MHz	3200MHz <sup>6</sup>
I/O Lanes	96 PCIe <sup>®</sup> 3.0	128 PCIe <sup>®</sup> 4.0 <sup>6</sup>
TDP	270Watts	200Watts
SPECrate <sup>®</sup> 2017_int_base	242	319
'Per Socket software' licensing cost	x2	x1
List Price	5800USD	4425USD




Cloud & Virtualization Performance - VMmark<sup>®</sup> 1.43X<sup>2</sup>




Floating-Point Performance - SPECrate<sup>®</sup>2017\_fp\_base<sup>2</sup>



JAVA<sup>®</sup> Performance - SPECjbb<sup>®</sup>2015-Multi-JVMmax-jOPS<sup>2</sup>




Integer Performance - SPECrate<sup>®</sup>2017\_int\_base<sup>2</sup>



HPC Performance - ANSYS Fluent<sup>®2</sup>





## AMD EPYC™ Processors

2<sup>nd</sup> and 3<sup>rd</sup> Gen AMD EPYC™ Processors have optimized multi-processor performance and scalability through the latest advancements in die-to-die connectivity with AMD Infinity Fabric™.

- High Speed Chiplet Interconnect
- Leadership Memory Bandwidth
  - 8 Memory Channels Per Socket
  - Peak 410 GB/s DRAM Bandwidth
- Leadership I/O Bandwidth
  - All 8 x16 links PCIe® Gen4
  - 128 PCIe® 4.0 Lanes in a Single Socket
  - 64GB/s bi-dir bandwidth per link, 512GB/s per socket



3rd Gen AMD EPYC™ server processors  
6-8 memory channel interleaving  
offering more configuration flexibility

Pause (k)

0:34 / 1:02

YouTube


The image shows a YouTube video player interface. The video content features a close-up of an AMD EPYC server processor mounted on a circuit board. The processor is a square chip with 'AMD EPYC' and 'D1C-1213' printed on it. The video player includes a title, a description, a progress bar, and standard playback controls like play/pause, volume, and settings. The background of the video is a dark, blue-tinted image of the server hardware.

**AMD** 3rd Gen AMD EPYC™ server processors

Watch later

Pause (k) **32MB of L3 cache per core**  
accelerating memory-intensive workloads

The image shows three server processors with glowing blue and purple outlines, set against a dark blue background with a grid pattern. A red underline is present under the text 'accelerating memory-intensive workloads'.



3rd Gen AMD EPYC™ server processors  
with new Zen3 cores  
delivering up to 19% IPC uplift<sup>2</sup>



Watch later

Sh

Pause (k)

**AMD** 3rd Gen AMD EPYC™ server processors

Watch later

Share

MORE VIDEOS

Scaling from 8 to 64 cores

# AMD Next Epyc

## AMD EPYC Genoa Pictured: 12 CPU Chiplets For 96 Zen 4 Cores

By [Anton Shilov](#) published 2 days ago

Since AMD's SP5 processor packaging is huge to enable 12 memory channels and higher thermal design power (up to 700W, according to some leaks), there is a lot of space for CPU chiplets and an I/O die. Indeed, we see 12 chiplets, each carrying eight Zen 4 cores. The flagship Genoa chip packs 96 cores, 192 threads, and AMD's next-generation I/O die supporting PCI Express 5.0 and several other innovations. However, AMD will also launch different SKUs with fewer cores; therefore, some of those chiplets will likely come disabled on some models. Another discovery on the chip is a rather complicated power delivery system with eight rows of capacitors.

Meanwhile, since the chiplets and the I/O die look too thick, we would say that the published image depicts a mockup of AMD's next-generation server processor rather than a sample of the product. Nonetheless, the photograph still shows AMD's Zen4-based EPYC 7004-series CPU without its integrated heat spreader.

# AMD Epyc: Genoa




## AMD EPYC Genoa Pictured: 12 CPU Chiplets For 96 Zen 4 Cores

By [Anton Shilov](#) published 2 days ago

Besides the transition from 64 cores to 96 cores, Genoa also wields the prowess of AMD's next-generation Zen 4 cores, PCIe 5.0 connectivity, and [DDR5](#) support. Since last year, AMD has started sampling Genoa to its partners, so the new 5nm data centers should debut this year.

2H2024



5th Gen AMD EPYC™ "Turin"

## 5th Gen AMD EPYC™ CPUs Available 2H 2024

"Zen 5" architecture

Up to 192 Cores and 384 threads

Confidential AI based on new Trusted IO feature

Leadership 4nm and 3nm process technologies

TSMC 4nm & 3nm

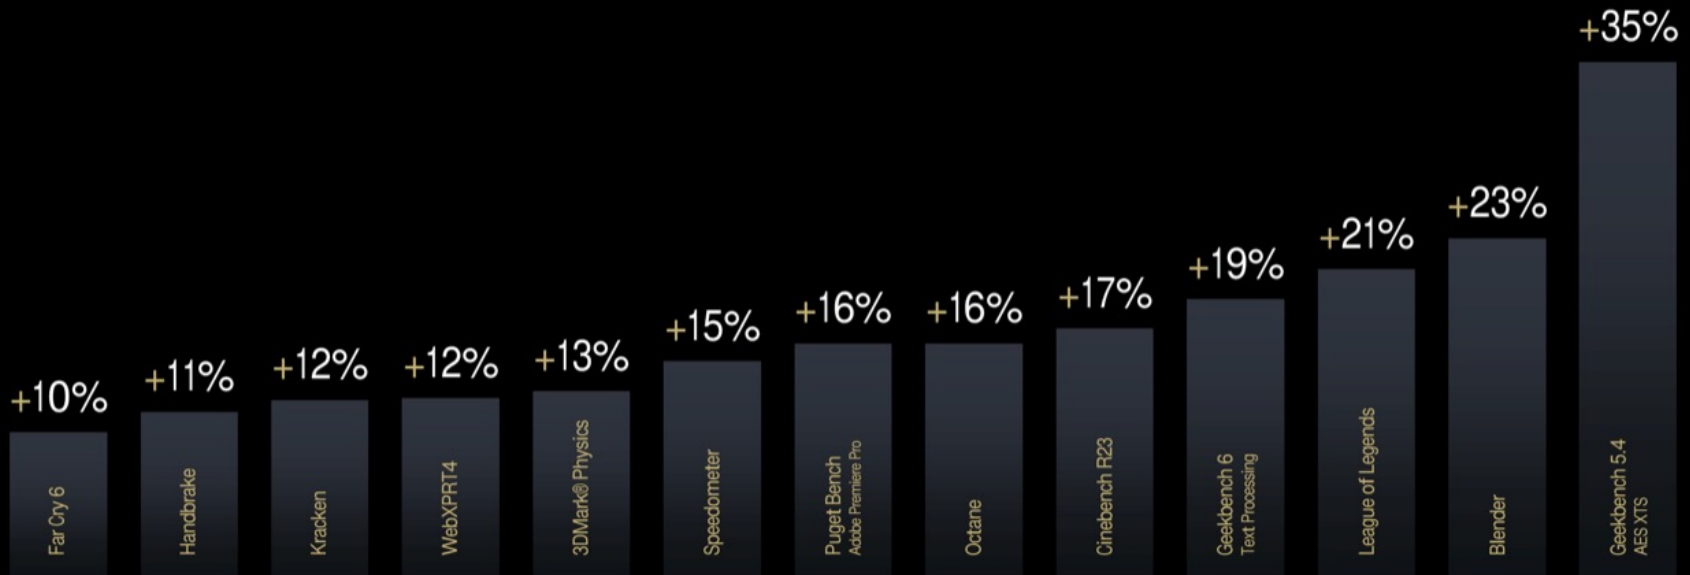
# AMD Zen 5

2H2024

## AMD Zen 5 Microarchitecture: 16% Better IPC Than Zen 4

Both the AMD Ryzen AI 300 series for mobile and the Ryzen 9000 series for desktops are powered by AMD's latest Zen 5 architecture, which brings a host of improvements in performance and efficiency. Perhaps the biggest improvement within their mobile lineup is the integration of the XDNA 2 NPU, which is designed to make use of the Microsoft Copilot+ AI software. These new mobile processors via the NPU can deliver up to 50 TOPS of AI performance, making them a significant upgrade in AMD's mobile chip lineup.

## "Zen5" 16% average IPC uplift\*





# AMD Zen 5



Integer Execution Advances

2H2024

## Wider Dispatch and Execute

- 8-wide dispatch/retire
- 6 ALU, 3 multiplies
- More unified ALU scheduler
- Large execution window





# AMD Zen 5

## Dual Pipe Fetch with Advanced Branch Prediction

- Branch prediction: less latency, more accuracy, and throughput
- Instruction cache latency and bandwidth improvement
- Dual decode pipes



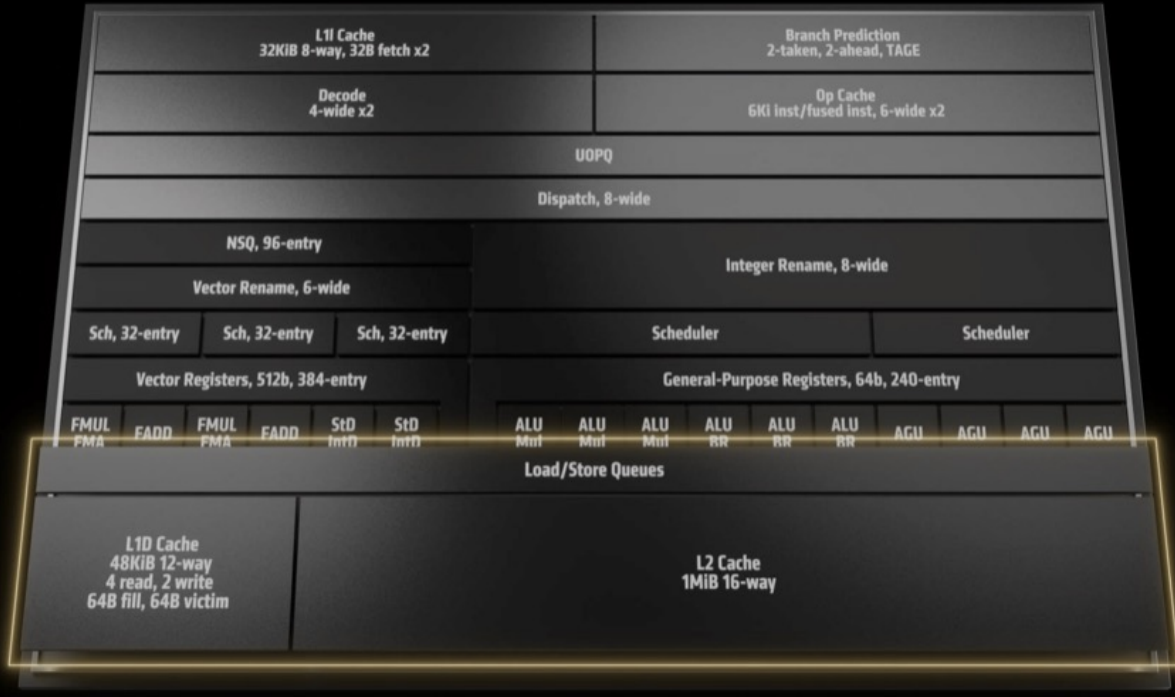
# AMD Zen 5

Load/Store Advances

2H2024

## Increased Data Bandwidth

- 48KB 12-way L1 data cache 4-cycle load
- Double the maximum bandwidth to the L1 cache and Floating-Point Unit
- Improved data prefetching



# AMD Zen 5



FP/Vector Math Unit Execution Advances

2H2024

## AVX-512

### 512-bit AI Datapath


- AVX-512 with full 512-bit data path
- 6 pipelines with two-cycle latency FADD
- Larger number of FP instructions in flight



# Section

## New HPC Products

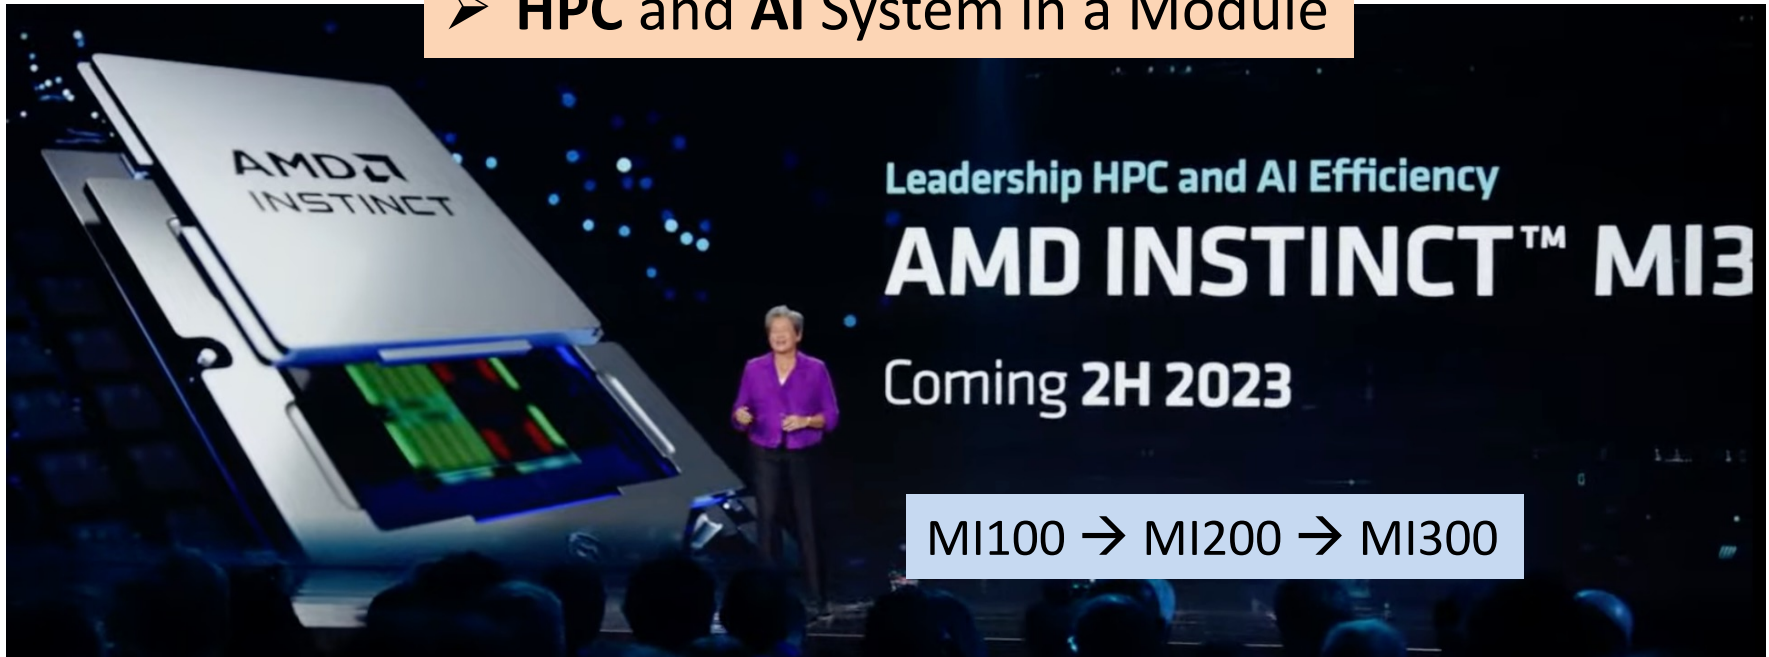
# Instinct MI300



# AMD CES 2023

Jan 2023

## ➤ HPC and AI System in a Module




Used in latest top **supercomputer** (Frontier)

# AMD CES 2023

Jan 2023

## AMD INSTINCT™ MI300

World's first data center integrated CPU + GPU



Next-Gen Accelerator Architecture



24 Leadership Data Center CPU cores

146B Transistors

128GB HBM3

3D

Advanced Chiplet Packaging in 5nm and 6nm process



AMD Instinct™ MI250X

AMD Instinct™ MI300



# Instinct MI300

## Supercharging AI and HPC

AMD Instinct™ MI300 Series accelerators are uniquely well-suited to power even the most demanding AI and HPC workloads, offering exceptional compute performance, large memory density, high bandwidth memory, and support for specialized data formats.

## AMD Instinct MI300X Accelerators

AMD Instinct MI300X Series accelerators are designed to deliver leadership performance for Generative AI workloads and HPC applications.





# Instinct MI300A

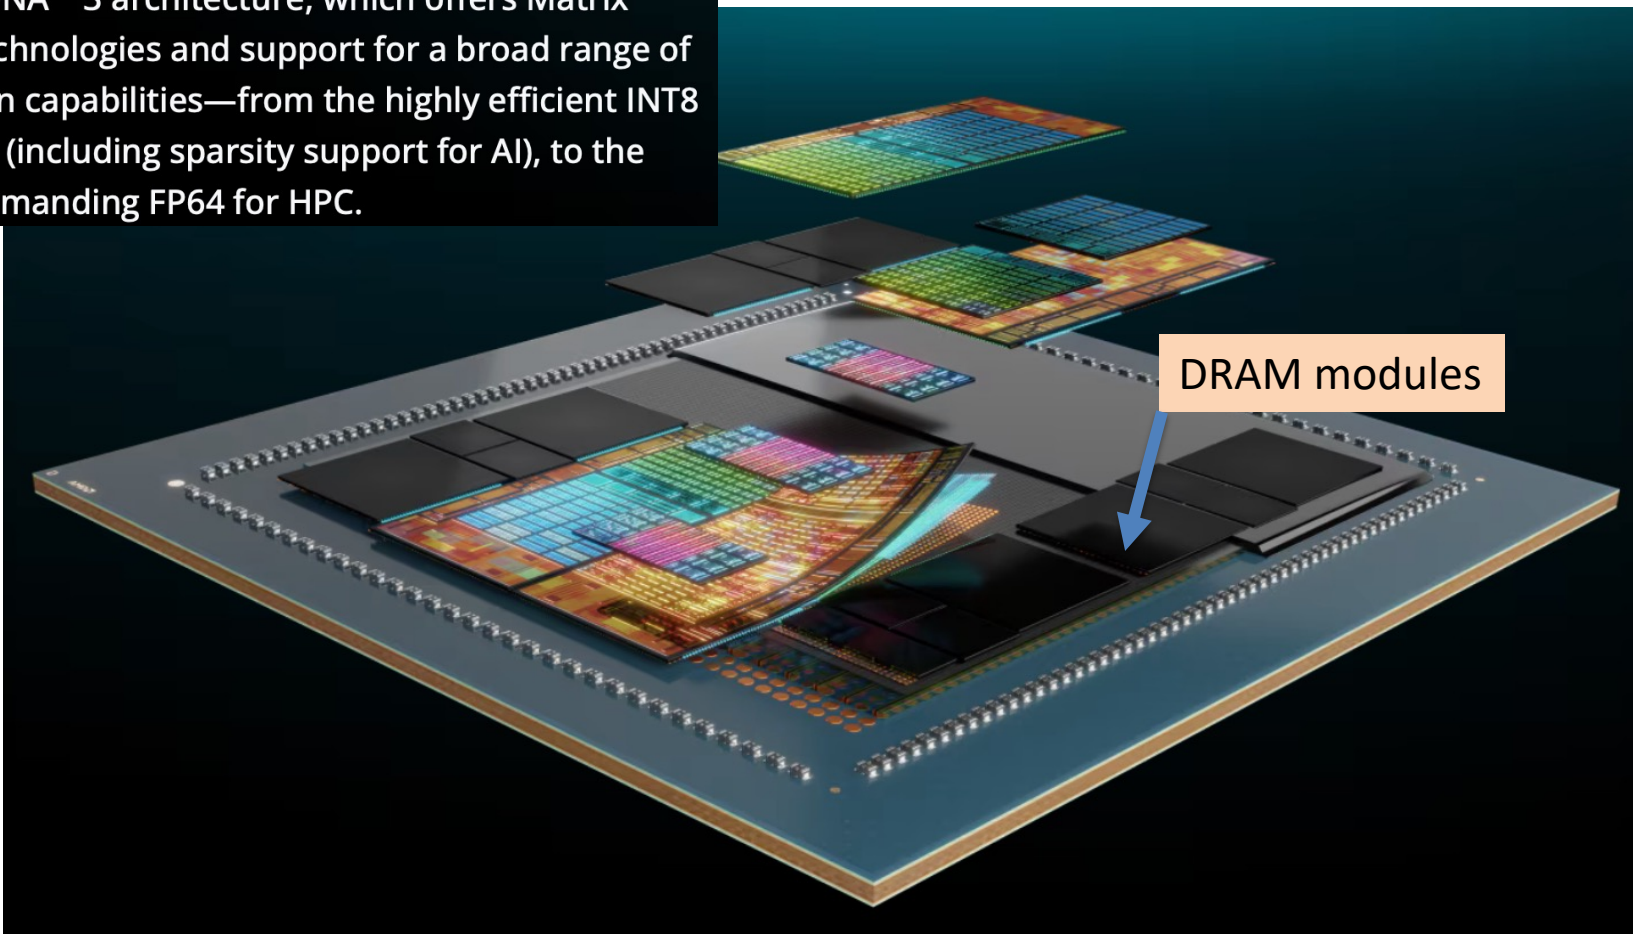


<b>Transistor Count</b>	153 Billion
-------------------------	-------------

# Instinct MI300A

**Under the Hood**

AMD Instinct MI300 Series accelerators are built on AMD CDNA™ 3 architecture, which offers Matrix Core Technologies and support for a broad range of precision capabilities—from the highly efficient INT8 and FP8 (including sparsity support for AI), to the most demanding FP64 for HPC.



3D Packaging: L3 V-Cache

# Instinct MI300A



Introducing AMD Instinct™ MI300 Series Accelerators  
**MI300A**  
Sharing 128 GB HBM3 Memory with 5.3 TB/s bandwidth

Share

Play (k)

0:16 / 1:38

YouTube

The video player shows a close-up of the MI300A accelerator chip with glowing blue lines representing data flow. The video title and key features are displayed at the top. The player controls include a play button, volume icon, progress bar, and a timestamp of 0:16 / 1:38. The YouTube logo and other interface icons are visible at the bottom.

**228 CUs**

228 GPU Compute Units

**24**

24 "Zen 4" x86 CPU Cores


**128 GB**

128 GB Unified HBM3  
Memory

**5.3 TB/s**

5.3 TB/s Peak Theoretical  
Memory Bandwidth

# Instinct MI300X



**304 CUs**

304 GPU Compute Units

**192 GB**

192 GB HBM3 Memory

**5.3 TB/s**

5.3 TB/s Peak Theoretical  
Memory Bandwidth

# Instinct MI300

COMP122 - **AMD** 

<b>GPU Architecture</b>	CDNA3		
<b>Lithography</b>	TSMC 5nm   6nm FinFET		
<b>Stream Processors</b>	19,456	<b>Peak Double Precision Matrix (FP64) Performance</b>	163.4 TFLOPs
<b>Matrix Cores</b>	1216	<b>Peak Single Precision (FP32) Performance</b>	163.4 TFLOPs
<b>Compute Units</b>	304		
<b>Peak Engine Clock</b>	2100 MHz	<b>Peak Double Precision (FP64) Performance</b>	81.7 TFLOPs
<b>Peak Eight-bit Precision (FP8)</b>	2.61 PFLOPs	<b>Peak INT8 Performance</b>	2.6 POPs

# MI300 Supercomputer



## Advancing Exascale Computing

AMD Instinct accelerators power some of the world's top supercomputers, including Lawrence Livermore National Laboratory's El Capitan system. See how this two-Exascale supercomputer will use AI to run first-of-its-kind simulations and advance scientific research.

# Section

---



# Other Features

# AMD Over-clocked

AMD Ryzen 9 vs. Intel i9

## AMD Update

### AMD News

#### **Overclocked Ryzen Sets New Performance Record.**

A liquid nitrogen-cooled AMD Ryzen 9 3900X processor has just set a new overclocking world record in the wPrime benchmark, beating the previous title holder, the Intel Core i9-7920X.

The feat was performed by Australian overclocker jordan.hyde99, who got the AMD Ryzen 9 3900X to reach speeds of 5,625 MHz to complete the tests in just 35 seconds and 517 milliseconds.


[Liquid nitrogen-cooled AMD Ryzen 9 3900X is the new world overclocking champ](#)

*Source: Tech radar (16 Dec 2019)*



## Modern Security Features with AMD Infinity Guard<sup>3</sup>

AMD EPYC™ Processors are designed with security in mind from the ground up to be highly resistant to today's sophisticated attacks. AMD Infinity Guard delivers a leading set of modern security features that help protect sensitive data, avoid downtime, and reduce resource drain.



# Section

---



# AMD Graphics

# AMD CPU's in PS5

7nm Ryzen



## PS5 Specs: Rumors and Current Predictions

Before we dive into our in-depth look at potential PS5 specs, let's start with the current predictions. The system's architect, Mark Cerny, revealed [some official details in April 2019](#), but the exact specs still remain a mystery. Here is what we know from that reveal:

- PS5 will support up to 8K resolutions
- The system will include an SSD that will drastically reduce load times
- Ray tracing (a powerful graphics technique) is supported by PS5
- The system uses a variation of AMD's third generation Ryzen with eight cores of the new 7nm Zen 2 microarchitecture
- GPU is a variation of the Radeon Navi family
- The system provides 3D audio without any additional hardware

# AMD CPU's in PS5



## 7nm Ryzen

These are our current spec predictions based on rumors and what was revealed by Mark Cerny:

- CPU: 8 core/16 threads at 3.2Ghz with a Zen2 architecture
- GPU: Navi-based with AMD next-gen features at 12.6 to 14.2 teraflops
- Memory: 24GB total with reportedly 20GB GDDR6 at 880GB/S and 4GB DDR4 for the operating system
- 2 TB SSD

# External GPU Box

## Nvidia GEFORCE RTX


The RTX 3090 will be the first 350W GPU form NVIDIA in quite a few years, possibly the biggest power hog ever. That to me does not seem like a step in the right direction. Part of that will be due to the also-very-hungry 24GB of GDDR6X. But the overall performance per watt ratio won't really be any better than the RTX 2080 Ti. A projected 39% boost in performance for a 40% boost in wattage.

It will probably be better than that, but not by a large amount.

I have long imagined a double sided GPU with cooling fans and heat sinks on both sides. The 3080/3090 FE will actually be double-sided to some degree as rumors point to a possible VRAM cooler mounted to cool the back of the PCB

350W

24GB GDDR6



# Section

---



# Quora Posts

# Quora Questions

---



## Related Questions

---

Why is AMD Ryzen beating Intel processors?

Which processor would you choose between AMD and Intel in 2020?

What is an AMD Ryzen processor?

Why is everyone preferring a AMD Ryzen than Intel Core processors?

Are AMD CPUs better for gaming than Intel CPUs?

Which has better processors for gaming purposes - AMD or Intel?



**Bob McConnell** · Sat

Former Vice President, Embedded Processors at Advanced Micro Devices  
(company) (1989–1999)

## Will Intel and AMD start making ARM processors to catch up with Apple?

Apple will not try to compete with Intel and AMD at the chip level. The PC market and the PC processor market has several things going on.

1. Apple has innovated and PC laptop makers like Dell, Lenovo, HP, and Acer will want some way to compete with Apple MacBooks. We have not yet seen the software that will soon come using the M1's neural engine.
2. Windows is a huge market and Microsoft has to decide how to proceed. Do they make a real serious attempt to create an ARM version? There are huge legacy issues, so they probably would have to make a second version of Windows to run on ARM. I think the laptop world is getting decided right now, but the desktop world may take a while.
3. Intel and AMD have built serious x86 processors for a long time. They'd love it if Microsoft stuck to x86 for Windows, but probably they will have to create ARM based processors. That's not a huge problem as ARM will license cores to anybody, but it'll open a new area of competition and a new area where things might go in different directions.

There is even room for some other company to jump into the processor market based on ARM and really stir the pot. But it won't be Apple.

194 views · View Upvoters · View Sharers · Answer requested by Kristian Fuler



You upvoted this



# My Response to Apple



**Jeff Drobman** · 2m ago

hey Bob, Apple has taken the lead in both process node at 5nm, giving them a new record of 16B transistors on the new M1 (even though the die is huge at 1"). they have innovated architecture for AI by adding 16 NPU cores along with extra ML accelerators. Intel and AMD have to decide how important AI is to their customer bases. this is way more significant than which ISA is used (IMHO).

**5-nanometer process**  
The first personal computer chip built with this cutting-edge technology.

**16 billion transistors**  
The most we've ever put into a single chip.

625 sq mm?  
(= 1 sq in)

# CPU – Serial

## The answer to the question

But why CPUs and GPUs are different in their parallelism:

Because their are made to run different code.

### CPU - serial

CPUs are made to run normal programs, written on normal programming languages, and to be able to handle ANY code.

Almost all common programming languages which are used to write normal programs are based on programming paradigms that are inherently serial - they consist of sequential operations. This stream of serial operations cannot be be efficiently parallelized in generic case, as the next operations very often need the results of the previous operations. Typically only very small parts of most programs can be automatically parallelized for multiple cores by the compiler. To parallelize most of the code for even few cores, the programmer has to do extra work for writing multi-threaded code (which is much harder, less productive and bug-prone).

So, CPUs are optimized for handling the serial code as quickly as possible. Because, for most programs, the is what really matters. If the parallel part is only half of the program, even if you have million cores and parallelize the code for those million cores, you are still spending half of the original execution time in the serial part, so you still are not even twice as fast as originally.

# GPU – Parallel

## GPUs - parallel


GPUs, on the other hand, are highly specialized for executing parallel workloads. They cannot even practically run whole programs. Only the small parts of the programs which are highly parallel can be sent to be executed in the GPU, and to do this, the programmer typically has to do lots of work to make his/hers code execute on the GPU.

Drawing the pixels to the screen is one of those very rare things which are very highly parallel - if the screen has 8 million pixels, and lots of calculations have to be done per pixel, theoretically colors of all the 8 millions pixels could be calculated in parallel.

Because of the highly parallel nature of graphics, these parallel processing units emerged as Graphics Processing Units(GPUs), even though for the last >10 years they have been also able to calculate parallel code that has nothing to do with graphics.


# Ryzen vs. i9: Power (TPD)

With the release of the Ryzen 5000 series, there is literally no reason to go for Intel's 10th gen Comet Lake chips (which still use the 2015 Skylake cores) given that AMD dominates single and multithreaded performance at every tier while their chips draw way less power and dissipate less heat.



# Epyc: *Server* Mkt Share

One motherboard with a pair of EPYC's is \$15,000. This is very big pill for a small IT department to swallow. Widespread adoption is going to be gradual, but you can already see in the graph you cited that AMD's share in the server market is its highest since 2011.



As I have discussed at length in recent posts, AMD architecture of the past decade has been horrible for efficiency, and thus entirely unfit for the server sector where efficiency and reliability are of the utmost importance.

EPYC 2 is extremely efficient and reliable. EPYC 3 will certainly be even more so. For AMD to make the leap that it has in 2020 is almost unprecedented in the history of computing. You might say that its installed base has leapt 500% in one quarter. At what time in history has any company ever increased its market share by 500% in three months? You might rightly say that its growth in 2020 is almost unprecedented in the entire history of commerce.

# Epyc for Servers

Customers like Twitter and Google are thrilled to have such powerful and compact options like the dual-EPYC (128-core, 256-thread 2U rackmount blades) that can in some cases replace an entire cabinet of old Xeon server blades while using a tiny fraction of the wattage.



# Section

---



# AMD APU's

# AMD APU's

40nm

2011

## Features overview [\[ edit \]](#)

The following table shows features of [AMD's APUs](#)

[\[ VisualEditor \]](#) [\[ view · talk · edit \]](#)

Codename	Server	Basic					Toronto				
		Micro									
	Desktop	Mainstream					Carrizo	Bristol Ridge	Raven Ridge	Picasso	Renoir
		Entry	Llano	Trinity	Richland	Kaveri					
		Basic									
	Mobile	Performance								Picasso	Renoir
		Mainstream	Llano	Trinity	Richland	Kaveri	Carrizo	Bristol Ridge	Raven Ridge		
		Entry									
		Basic									
	Embedded			Trinity			Bald Eagle	Merlin Falcon, Brown Falcon			Great Horned Owl
	Platform	High, standard and low power									
	Released	Aug 2011	Oct 2012	Jun 2013	Jan 2014	Jun 2015	Jun 2016	Oct 2017	Jan 2019	Mar 2020	
CPU microarchitecture	K10	Piledriver		Steamroller	Excavator	"Excavator+" <sup>[1]</sup>	Zen	Zen+	Zen 2		
ISA	x86-64										



# AMD APU's

COMP122

40nm

2011

Embedded APUs [\[ edit \]](#)

Further information: [Embedded system](#)

**G-Series** [\[ edit \]](#)

**Brazos: "Ontario" and "Zacate" (2011)** [\[ edit \]](#)

- [Fabrication](#) 40 nm
- [Socket FT1](#) (BGA-413)
- CPU microarchitecture: [Bobcat](#)<sup>[164]</sup>
- L1 Cache: 32 KB Data per core and 32 KB Instructions per core
- [MMX](#), [SSE](#), [SSE2](#), [SSE3](#), [SSSE3](#), [SSE4a](#), [ABM](#), [NX bit](#), [AMD64](#), [AMD-V](#)
- GPU microarchitecture: [TeraScale 2 \(VLIW5\)](#) "Evergreen"
- Memory support: single-channel, support up to two DIMMs of DDR3-1333 or DDR3L-1066
- 5 GT/s UMI

Model ↕	Released ↕	Fab ↕	Stepping ↕	CPU			GPU				Memory support ↕	TDP ↕		
				Cores (threads)/[FPUs] ↕	Clock ↕	Cache <sup>[a]</sup>			Model ↕	Config ↕			Clock ↕	Processing power (GFLOPS) <sup>[b]</sup> ↕
						L1 ↕	L2 ↕	L3 ↕						
<b>G-Series T24L</b>	Mar 1, 2011 May 23, 2011		B0	1 (1) [1]	800 MHz	512 KB					DDR3-1066	5 W		
<b>G-Series T30L</b>	Mar 1, 2011 May 23, 2011				1.0 GHz							18 W		
<b>G-Series T48L</b>	Mar 1, 2011 May 23, 2011			2 (2) [1]	1.4 GHz	2 × 512 KB								

# AMD APU's

14-7nm 2011

Fab. (nm)	GF 32SHP (HKMG SOI)		GF 28SHP (HKMG bulk)			GF 14LPP (FinFET bulk)	GF 12LP (FinFET bulk)	TSMC N7 (FinFET bulk)	
Die area (mm <sup>2</sup> )	228	246	245	245	250	210 <sup>[3]</sup>		156	
Min TDP (W)	35	17		12			10		
Max APU TDP (W)	100		95	65					
Max stock APU base clock (GHz)	3	3.8	4.1	4.1	3.7	3.8	3.6	3.7	3.8
Max APUs per node <sup>[b]</sup>	1								
Max CPU <sup>[c]</sup> cores per APU	4							8	
Max threads per CPU core	1						2		
Integer structure	3+3	2+2				4+2		4+2+1	
I386, i486, i586, CMOV, NOPL, i686, PAE, NX bit, CMPXCHG16B, AMD-V, RVI, ABM, and 64-bit LAHF/SAHF	✓								
IOMMU <sup>[d]</sup>	N/A	✓							
BMI1, AES-NI, CLMUL, and F16C	✓								
MOVBE	N/A		✓						
AVIC, BMI2 and RDRAND	✓								
ADX, SHA, RDSEED, SMAP, SMEP, XSAVEC, XSAVES, XRSTORS, CLFLUSHOPT, and CLZERO	N/A					✓			
WBNOINVD, CLWB, RDPID, RDPRU, and MCOMMIT	N/A							✓	
FPU's per core	1	0.5				1			
Pipes per FPU	2								
FPU pipe width	128-bit							256-bit	
CPU instruction set SIMD level	SSE4a <sup>[e]</sup>	AVX			AVX2				

# AMD APU's

As of May 1, 2013, AMD opened the doors of their "semi-custom" business unit.<sup>[175]</sup> Since these chips are custom-made for specific customer needs, they vary widely from both consumer-grade APUs and even the other custom-built ones. Some notable examples of semi-custom chips that have come from this sector include the chips from the [PlayStation 4](#) and [Xbox One](#).<sup>[176]</sup> So far the size of the integrated GPU in these semi-custom APUs exceed by far the GPU size in the consumer-grade APUs.

Chip (device) ↕	Release date ↕	Fab ↕	Die area (mm2) ↕	CPU				GPU			
				Architecture ↕	Cores ↕	Clock (GHz) ↕	L2 cache ↕	Architecture ↕	Core config <sup>[a]</sup> ↕	Clock (MHz) ↕	GFLOPS <sup>[b]</sup> ↕
<b>Liverpool (PS4)</b>	Nov 2013	28 nm	348	Jaguar	2 modules with 4 cores	1.6	2x 2 MiB	GCN 2	1152:72:32 18 CU	800	1843
<b>Durango (Xbox One)</b>	Nov 2013		363						853	1310	
<b>Edmonton (Xbox One S)<sup>[177]</sup></b>	Jun 2016	240				1.75			768:48:16 12 CU	914	1404

# AMD APU's

COMP122

28nm

2016

I-Family: "Brown Falcon" (2016, SoC) [edit]

- Fabrication 28 nm
- Socket FP4<sup>[166]</sup>
- 2 or 4 Excavator x86 cores with 1MB shared L2 cache
- L1 Cache: 32 KB Data per core and 96 KB Instructions per module
- MMX, SSE, SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, SSE4a, AMD64, AMD-V, AES, CLMUL, AVX, AVX 1.1, AVX2, XOP, FMA3, FMA4, F16C, ABM, BMI1, BMI2, TBM, RDRAND
- GPU microarchitecture: Graphics Core Next (GCN) (up to 4 CUs) with support for DirectX 12
- Dual channel 64-bit DDR4 or DDR3 memory with ECC
- 4K x 2K H.265 decode capability and multi format encode and decode
- Integrated Controller Hub supports: PCIe 3.0 1x4, PCIe 2/3 4x1, 2 USB3 + 2 USB2 ports, 2 SATA 2.0/3.0 ports

Model	Released	Fab	Stepping	CPU						GPU					Memory support
				[Modules/FPUs] Cores/threads	Clock	Turbo	Cache <sup>[a]</sup>			Model	Config <sup>2</sup>	Clock	Turbo	Processing power (GFLOPS) <sup>[b]</sup>	
							L1	L2	L3						
GX-217GI	Feb 23, 2016	28nm		[1] 2	1.7 GHz	2.0 GHz	96 KB inst. per module	1 MB	N/A	R6E	256:16:4 4 CU	758 MHz	N/A	388	DDR3/DDR4-1600
GX-420GI <sup>[167][168]</sup>	2016			[2] 4	2.0 GHz	2.2 GHz	32 KB data per core	2 MB		R6E	256:16:4 4 CU	758 MHz	N/A	388	DDR4-1866
									R7E	384:24:4 6 CU	626 MHz		480.7		

# AMD SoC's

COMP122

28nm

2016-17

LX-Family (2016, SoC) [edit]

- Fabrication 28 nm
- Socket **FT3b** (769-BGA)
- 2 **Puma** x86 cores with 1MB shared L2 cache
- L1 Cache: 32 KB Data per core and 32 KB Instructions per core
- **MMX, SSE, SSE2, SSE3, SSSE3, SSE4a, SSE4.1, SSE4.2, AVX, F16C**, CLMUL, **AES**, MOVBE (Move Big-Endian instruction), XSAVE/XSAVEOPT, **ABM, BMI1, BMI2**
- GPU microarchitecture: **Graphics Core Next** (GCN) (1CU) with support for DirectX 11.2
- Single channel 64-bit DDR3 memory with **ECC**
- Integrated Controller Hub supports: PCIe® 2.0 4x1, 2 USB3 + 4 USB2 ports, 2 SATA 2.0/3.0 ports

Model	Released ↕	Fab ↕	Stepping ↕	CPU					GPU				Memory support ↕	TDP ↕			
				Cores (threads) ↕ [FPUs]	Clock ↕	Cache <sup>[a]</sup>			Model ↕	Config ↕	Clock ↕	Processing power (GFLOPS) <sup>[b]</sup> ↕					
						L1 ↕	L2 ↕	L3 ↕									
<b>GX-208JL</b>	Feb 23, 2016	28nm	ML-A1	2	800 MHz	32 KB inst. 32 KB data	1 MB	N/A	R1E	64:4:1 1 CU	267 MHz	34.1	DDR3-1333	6 W			
<b>GX-210HL</b>	2017												DDR3-1066	7 W			
<b>GX-210JL</b>	Feb 23, 2016				1.0 GHz									DDR3-1333	6 W		
<b>GX-210KL</b>	2017													4.5 W			
<b>GX-215GL</b>	Feb 23, 2016				1.5 GHz								per core	497 MHz	63.6	DDR3-1600	15 W
<b>GX-218GL</b>					1.8 GHz												

# AMD SoC's

COMP122

28nm

2017

## Opteron X3000-series "Toronto" (2017) [\[edit\]](#)

- Fabrication 28 nm
- Socket **FP4**
- Two or Four CPU cores based on the **Excavator** microarchitecture<sup>[93][94]</sup>
- L1 Cache: 32 KB Data per core and 96 KB Instructions per module
- MMX, SSE, SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, SSE4a, AMD64, AMD-V, AES, CLMUL, AVX, AVX 1.1, AVX2, XOP, FMA3, FMA4, F16C, ABM, BMI1, BMI2
- **DDR4 SDRAM**
- GPU based on **3rd Generation GCN (Graphics Core Next)** architecture

Model	Released	Fab	Stepping	CPU						GPU				Memory support	TDP
				Cores (threads)	Clock	Turbo	Cache <sup>[a]</sup>			Model	Config	Clock	Processing power (GFLOPS) <sup>[b]</sup>		
							L1	L2	L3						
X3216	Q2 2017	28nm	01h	4 (4)	1.6 GHz	3.0 GHz	96 KB inst. per module	1 MB	N/A	R5	256:16:4 4 CU	800 MHz	409.6	DDR4-1600	12-15 W
X3418				4 (4)	1.8 GHz	3.2 GHz		2 MB		R7	384:24:6 6 CU		614.4		
X3421				4 (4)	2.1 GHz	3.4 GHz	32 KB data per core	2 MB		R7	512:32:8 8 CU		819.2		

# AMD SoC's

COMP122

28nm

2017-18

J-Family: "Prairie Falcon" (2016, SoC) [edit]

- Fabrication 28 nm
- Socket FP4<sup>[169]</sup>
- 2 "Excavator+" x86 cores with 1MB shared L2 cache
- L1 Cache: 32 KB Data per core and 96 KB Instructions per module
- MMX, SSE, SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, SSE4a, AMD64, AMD-V, AES, CLMUL, AVX, AVX 1.1, AVX2, XOP, FMA3, FMA4, F16C, ABM, BMI1, BMI2, TBM, RDRAND
- GPU microarchitecture: Radeon R5E Graphics Core Next (GCN) (up to 3 CUs) with support for DirectX 12
- Single channel 64-bit DDR4 or DDR3 memory
- 4K x 2K H.265 decode capability with 10-bit compatibility and multi format encode and decode
- Integrated Controller Hub supports: PCIe 3.0 1x4, PCIe 2/3 4x1, 2 USB3 + 2 USB2 ports, 2 SATA 2.0/3.0 ports

Model	Released ↕	Fab ↕	Stepping ↕	CPU						GPU					Memory support ↕	TDP ↕
				[Modules/FPUs] Cores/threads ↕	Clock ↕	Turbo ↕	Cache <sup>[a]</sup>			Model ↕	Config <sup>2</sup> ↕	Clock ↕	Turbo ↕	Processing power (GFLOPS) <sup>[b]</sup> ↕		
							L1 ↕	L2 ↕	L3 ↕							
GX-212JJ	2018	28nm		[1] 2	1.2 GHz	1.6 GHz	96 KB inst. per module	1 MB	N/A	R1E	64:4:1 1 CU	600 MHz	N/A	76.8	DDR3-1333 DDR4-1600	6–10 W
GX-215JJ	2017				1.5 GHz	2.0 GHz				R2E	128:8:2 2 CU			153.6	DDR3-1600 DDR4-1866	
GX-220IJ	2018				2.0 GHz	2.2 GHz	32 KB data per core	R4E	192:12:3 3 CU	230.4	DDR3-1866 DDR4-2133					
GX-224IJ	2017				2.4 GHz	2.8 GHz										

# AMD SoC's

COMP122

14nm

2018

**1000-Series** [edit]

**V1000-Family: "Great Horned Owl" (2018, SoC)** [edit]

- Fabrication 14 nm by GlobalFoundries
- Up to 4 Zen cores
- Socket FP5
- MMX, SSE, SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, SSE4a, AMD64, AMD-V, AES, CLMUL, AVX, AVX 1.1, AVX2, FMA3, F16C, ABM, BMI1, BMI2, RDRAND, Turbo Core
- Dual channel DDR4 memory with ECC
- Fifth generation GCN based GPU

Model	Release date	Process	CPU						GPU				Memory support	Ethernet	TDP (W)	Junction temperature (°C)
			Cores (threads)	Clock rate (GHz)		Cache <sup>[i]</sup>			Model	Config <sup>[ii]</sup>	Clock	Processing power (GFLOPS) <sup>[iii]</sup>				
				Base	Boost	L1	L2	L3								
<b>V1500B</b> <sup>[172]</sup>	December 2018	GloFo 14LP	4 (8)	2.2	N/A	64 KB inst. 32 KB data per core	512 KB per core	4 MB	N/A				DDR4-2400 dual-channel	2x 10GbE	12–25	0–105
<b>V1780B</b> <sup>[172]</sup>				3.35	3.6								DDR4-3200 dual-channel		35–54	
<b>V1202B</b> <sup>[172]</sup>	February 2018		2 (4)	2.3	3.2				RX Vega 3	192:12:16 3 CU	1000 MHz	384	DDR4-2400 dual-channel		12–25	
<b>V1404</b> <sup>[172]</sup>	December 2018			2.0					RX Vega	512:32:16	1100 MHz	1126.4			–40 – 105	



# AMD

COMP122

12nm

2019

## "Picasso" (2019) [\[ edit \]](#)

Main article: [Ryzen](#)

- Fabrication 12 nm by GlobalFoundries
- Socket FP5
- Up to four Zen+ CPU cores
- MMX, SSE, SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, SSE4a, AMD64, AMD-V, AES, CLMUL, AVX, AVX 1.1, AVX2, FMA3, F16C, ABM, BMI1, BMI2, RDRAND, Turbo
- Dual-channel DDR4 memory controller
- Fifth generation GCN based GPU

Model	Release date	Process	CPU						GPU				Memory support	TDP
			Cores (threads)	Clock rate (GHz)		Cache <sup>[i]</sup>			Model	Config <sup>[ii]</sup>	Clock	Processing power (GFLOPS) <sup>[iii]</sup>		
				Base	Boost	L1	L2	L3						
Ryzen 3 3300U <sup>[116]</sup>	January 6, 2019		4 (4)	3.5				Vega 6	384:24:8 6 CU <sup>[117]</sup>	1200 MHz	921.6		15 W	
Ryzen 3 PRO 3300U <sup>[118]</sup>														
Ryzen 5 3500U <sup>[119]</sup>														
Ryzen 5 PRO 3500U <sup>[121]</sup>	September 22,			2.1					512:32:16 8 CU <sup>[120]</sup>		1228.8			
Ryzen 5														

# AMD

## COMP122

14nm

2019-20

**R1000-Family: "Banded Kestrel" (2019, SoC)** [ edit ]

- Fabrication 14 nm by GlobalFoundries
- Up to 2 Zen cores
- Socket FP5
- MMX, SSE, SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, SSE4a, AMD64, AMD-V, AES, CLMUL, AVX, AVX 1.1, AVX2, FMA3, F16C, ABM, BMI1, BMI2, RDRAND, Turbo Core
- Dual channel DDR4 memory with ECC
- Fifth generation GCN based GPU

Model	Release date	Process	CPU							GPU				Memory support	TDP
			Cores (threads)	Clock rate (GHz)			Cache <sup>[i]</sup>			Model	Config <sup>[ii]</sup>	Clock	Processing power (GFLOPS) <sup>[iii]</sup>		
				Base	Boost	XFR	L1	L2	L3						
<b>R1102G</b> <sup>[174]</sup>	February 25, 2020	GloFo 14LP	2 (2)	1.2	2.6	Unknown	64KB inst. 32 KB data per core	512KB per core	4MB	RX Vega 3	192:12:4 3 CU	1000 MHz	384	DDR4-2400 single-channel	6 W
<b>R1305G</b> <sup>[174]</sup>			1.5	2.8	Unknown										
<b>R1505G</b> <sup>[174]</sup>	April 16, 2019		2 (4)	2.4	3.3	Unknown						1200 MHz	460.8	DDR4-2400 dual-channel	12–25 W
<b>R1606G</b> <sup>[174]</sup>			2.6	3.5	Unknown										

## Xbox

<b>Scarlett (Xbox Series X)</b>	2020 planned	360.45	3.66 (3.8 w/o SMT)	3328:?:? 52 CU	1825	12155	Double-rate FP16 Real-time ray tracing Mesh shaders Variable rate shading ANN acceleration	10 GiB 6 GiB	GDDR6 320-bit GDDR6 192-bit <sup>[i]</sup>	560 336
---------------------------------	--------------	--------	--------------------	----------------	------	-------	--	-----------------	---	------------

# AMD

COMP122

14nm

2020

## "Dali", "Pollock" (2020) [\[ edit \]](#)

- **Fabrication** 14 nm by [GlobalFoundries](#)
- Socket FP5 / FT5
- Two [Zen](#) CPU cores
- MMX, [SSE](#), [SSE2](#), [SSE3](#), [SSSE3](#), [SSE4.1](#), [SSE4.2](#), [SSE4a](#), [AMD64](#), [AMD-V](#), [AES](#), [CLMUL](#), [AVX](#), [AVX 1.1](#), [AVX2](#), [FMA3](#), [F16C](#), [ABM](#), [BMI1](#), [BMI2](#), [RDRAND](#), [T](#)

Model ↕	Release date ↕	Fab ↕	CPU						GPU				Memory support ↕
			Cores/(threads) ↕	Clock rate (GHz)		Cache <sup>[a]</sup>			Model ↕	Config <sup>[b]</sup> ↕	Clock ↕	Processing power (GFLOPS) <sup>[c]</sup> ↕	
				Base ↕	Boost ↕	L1 ↕	L2 ↕	L3 ↕					
<b>AMD 3015e</b> <sup>[155]</sup>	July 6, 2020		2 (4)	1.2	2.3					3 CU	600 MHz		DDR4-1600 single-channel
<b>AMD 3020e</b> <sup>[156]</sup>			2 (2)	1.2	2.6					192:12:4	1000 MHz	384	
<b>Athlon Silver</b>	January 6,		2 (4)	1.4	2.8					3 CU	MHz		

7nm

2020

**"Renoir" (2020)** [edit]

Main article: [Ryzen](#)

- Fabrication 7 nm by TSMC<sup>[131][132][133]</sup>
- Socket FP6
- 9.8 billion transistors on one single 7nm monolithic die<sup>[134]</sup>
- Up to eight [Zen 2](#) CPU cores

Model	Release date	Process	CPU						GPU			Memory support	TDP										
			Cores (threads)	Clock rate (GHz)		Cache <sup>[a]</sup>			Model & config <sup>[b]</sup>	Clock	Processing power (GFLOPS) <sup>[c]</sup>												
				Base	Boost	L1	L2	L3															
<b>Ryzen 3 4300U</b> <sup>[135]</sup>	March 16, 2020		4 (4)	2.7	3.7			4 MB	7nm Vega 5	1400 MHz	896		10-25W										
<b>Ryzen 3 PRO 4450U</b> <sup>[136]</sup>	May 7, 2020		4 (8)	2.5					320:20:8														
<b>Ryzen 5 4500U</b> <sup>[137]</sup>	March 16, 2020		6 (6)	2.3	4.0										7nm Vega 6	1500 MHz	1152		35 W				
<b>Ryzen 5 4600U</b> <sup>[138]</sup>			6 (12)	2.1																384:24:8			
<b>Ryzen 5 PRO 4650U</b> <sup>[139]</sup>	May 7, 2020		6 (12)																				6 CU
<b>Ryzen 5 4600HS</b> <sup>[140]</sup>																							

# AMD GPU Timeline

COMP122

				150-7nm		2020				
<b>R200</b>	Programmable pixel & vertex pipelines	150nm	N/A	2.0 <sup>[b]</sup>	8.1	N/A	Ended	2001		
<b>R300</b>		150nm 130nm 110nm			9.0 11 (FL 9_2)			2002		
<b>R420</b>		130nm 110nm			9.0b 11 (FL 9_2)			2004		
<b>R520</b>		90nm 80nm			9.0c 11 (FL 9_3)			2005		
<b>R600</b>	TeraScale 1	80nm 65nm		3.3	10.0 11 (FL 10_0)			ATI Stream	2007	
<b>RV670</b>		55nm			10.1 11 (FL 10_1)			ATI Stream APP <sup>[17]</sup>	2007	
<b>RV770</b>		55nm 40nm			1.0			2008		
<b>Evergreen</b>	TeraScale 2	40nm		4.5 (Linux 4.2) <sup>[18][19][20][c]</sup>	11 (FL 11_0)			1.2	2009	
<b>Northern Islands</b>	TeraScale 2 TeraScale 3								2010	
<b>Southern Islands</b>	GCN 1 <sup>st</sup> gen	28nm		1.0	4.6 (Mesa 4.5)			✓	1.2 2.0 possible	2012
<b>Sea Islands</b>	GCN 2 <sup>nd</sup> gen		1.1			2.0 (1.2 in MacOS, Linux) 2.1 Beta in Linux ROCm 2.2 possible	Current		2013	
<b>Volcanic Islands</b>	GCN 3 <sup>rd</sup> gen								2014	
<b>Arctic Islands</b>	GCN 4 <sup>th</sup> gen								14nm	2016
<b>Vega</b>	GCN 5 <sup>th</sup> gen								14nm 7nm	11 (FL 12_1) 12 (FL 12_1)
<b>Navi</b>	RDNA 1 <sup>st</sup> gen		7nm				2019			

# AMD GPU

>40nm 2020

## Graphics API overview [\[ edit \]](#)

The following table shows the graphics and compute APIs support across AMD GPU microarchitectures. Note that a branding series might include

[\[ VisualEditor \]](#) [\[ view · talk · edit \]](#)

Chip series	Micro-architecture	Fab	Supported APIs					AMD support	Year introduced
			Rendering			Computing			
			Vulkan <sup>[15]</sup>	OpenGL <sup>[16]</sup>	Direct3D	HSA	OpenCL		
<b>Wonder</b>	Fixed-pipeline <sup>[a]</sup>	1000nm	N/A	N/A	N/A	N/A	N/A	1986	
<b>Mach</b>		800nm						1991	
<b>3D Rage</b>		600nm						1996	
<b>Rage Pro</b>		500nm						1997	
<b>Rage 128</b>		350nm						1998	
<b>R100</b>		250nm						2000	
<b>R200</b>	Programmable pixel & vertex pipelines	180nm	N/A	2.0 <sup>[b]</sup>	N/A	N/A	Ended	2001	
<b>R100</b>		150nm						2000	
<b>R200</b>		150nm						2001	
<b>R300</b>		130nm						2002	
<b>R420</b>		110nm						2004	
<b>R520</b>		90nm						2005	
<b>R600</b>	TeraScale 1	80nm	N/A	3.3	N/A	N/A	Ended	2007	
<b>R600</b>		65nm						2007	
<b>RV670</b>		55nm						2007	
<b>RV770</b>		55nm						2008	
<b>Evergreen</b>	TeraScale 2	40nm	4.5	11 (FL 10_1)	1.0			2009	

# AMD GPU

>14nm

2016-20

L1 data cache per core (KiB)	64	16		32		64		128		
L1 data cache associativity (ways)	2	4		8		16		32		
L1 instruction caches per core	1	0.5		1		2		4		
Max APU total L1 instruction cache (KiB)	256	128	192		256	512		1024		
L1 instruction cache associativity (ways)	2		3		4	8		16		
L2 caches per core	1	0.5		1		2		4		
Max APU total L2 cache (MiB)	4		8		16		32		64	
L2 cache associativity (ways)	16		32		64		128		256	
APU total L3 cache (MiB)	N/A		N/A		4		8		16	
APU L3 cache associativity (ways)	N/A		N/A		16		32		64	
L3 cache scheme	Victim		N/A		Victim		Victim		Victim	
Max stock DRAM support	DDR3-1866	DDR3-2133		DDR3-2133, DDR4-2400	DDR4-2400	DDR4-2933		DDR4-3200, LPDDR4-4266		
Max DRAM channels per APU	2		2		2		2		2	
Max stock DRAM bandwidth (GB/s) per APU	29.866	34.132		38.400	46.932		68.256		102.400	
GPU microarchitecture	TeraScale 2 (VLIW5)	TeraScale 3 (VLIW4)		GCN 2nd gen	GCN 3rd gen		GCN 5th gen <sup>[4]</sup>			
GPU instruction set	TeraScale instruction set			GCN instruction set						
Max stock GPU base clock (MHz)	600	800	844	866	1108		1250	1400	2100	
Max stock GPU base GFLOPS <sup>[1]</sup>	480	614.4	648.1	886.7	1134.5		1760	1971.2	2150.4	
3D engine <sup>[9]</sup>	Up to 400:20:8	Up to 384:24:6		Up to 512:32:8			Up to 704:44:16 <sup>[5]</sup>		Up to 512:?:?	
	IOMMUv1			IOMMUv2						

# Section

---



## Software Tools

<https://www.amd.com/en/technologies>

<https://www.amd.com/en/technologies/ryzen-master>



## Tools & SDKs

### AMD Zen Software Studio

**AMD Optimizing C/C++ and Fortran Compilers (“AOCC”)** — The AOCC compiler system is a high performance, production software generation tool optimized for AMD processors based on the AMD “Zen” core architecture.

**AMD  $\mu$ Prof** — AMD  $\mu$ Prof is a suite of powerful tools that help developers optimize software for performance and power, optimized for AMD processors based on the AMD “Zen” core architecture.

**AMD Optimizing CPU Libraries (“AOCL”)** — AOCL is a set of numerical libraries optimized for AMD processors based on the AMD “Zen” core architecture.

### Other SDKs and Tools

**Tools for DMTF DASH** — DASH (Desktop and mobile Architecture for System Hardware) is a client management standard released by the DMTF (Distributed Management Task Force). DASH is a web services based standard for secure out-of-band and remote management of desktops and mobile systems. Client systems that support out-of-band management help IT administrators perform tasks independent of the power state of the machine or the state of the operating system.

**AMD Ryzen™ Master Monitoring SDK** — The AMD Ryzen™ Master Monitoring SDK is a public distribution

# AMD Software Tools



### HIPIFIED (PORTED) CODES

QUICKSILVER	HACC
SWALITE	HPL
PENNANT	LAGHOS

AND (ALREADY) MANY MORE...

**AMD**

### AMD GPU COMPILERS

- AMD supports several compilers that emit AMDGCN assembly
  - hcc
  - ADMP (in development, but reworking for completeness)
  - To inspect AMDGCN, instead of objdump you need xct.cact.hccmcc1 (installed with hcc)
  - The GCN ISA is free and open: <https://developer.amd.com/resources/Downloads/gcn-isa.html>
- HCC
  - Invoked by hipcc and is just a Perl script that wraps around hcc
  - Compiles HIP code
  - HIP (Heterogeneous Interface for Portability) is an interface that looks similar to CUDA
  - hcc is a fork of clang
    - understands HIP and emits AMDGCN in the resulting binary
      - o bin/clang-obj-amdgcn
      - % pieces are dealt with in the same way
      - OpenMP Compiler
      - code with OpenMP "target" pragmas
      - right to produce a binary that can offload work to the GPU

```

graph TD
    hipcc --> hcc
    hcc --> amdgcn
    
```

**AMD**

### SIMD EXECUTION

```

// (threadIdx.x * 2) {
    #pragma omp for
}
// (threadIdx.x * 2) {
    //}
//} else {
    //}
//}
    
```

**AMD**

## Porting CUDA to HIP

In the final video of the series, presenter Nicholas Malaya...

## GPU Programming Software

In this video, presenter Damon McDougall summarizes the various Compilers,...

## GPU Programming Concepts (Part 3)

In this video, presenter Noel Chalmers concludes the discussion on...

# Software Tools



## AMD EPYC™ 7xx2-series Processors Compiler Options Quick Reference Guide

### **AOCC compiler (with Flang - Fortran Front-End)**

Latest release: 2.1, Nov 2019

<https://developer.amd.com/amd-aocc/>

#### **Architecture**

Generate instructions that run on 2nd Gen EPYC/ RYZEN	-march=znver2
---	---------------

Generate instructions for the local machine	-march=native
---	---------------

#### **Optimization Levels**

Disable all optimizations	-O0
---------------------------	-----

Minimal level speed and code optimization	-O1
---	-----

Moderate level optimization (default)	-O2/ -O
---------------------------------------	---------

Aggressive optimizations	-O3
--------------------------	-----

Maximize performance	-Ofast
----------------------	--------

# AMD Optimizing C/C++ Compiler

**NEW! AOCC 3.0 is now available**

[📄 Downloads](#) [📄 User Guide](#)

The AOCC system is a high performance, production quality code generation tool. The AOCC environment C++ and Fortran applications targeting 32-bit and 64-bit Linux<sup>®</sup> platforms. The AOCC system offers advanced global optimization, vectorization, inter-procedural analyses, loop transformations and code generation. AM from each x86 processor core when utilized. The AOCC suite simplifies and accelerates the development a

AOCC offers:

- Enabled, tuned and optimized for AMD EPYC architectures
- Based on LLVM Compiler Infrastructure
- Tuned for AMDLibM (AMD Math library)
- Flang as the default Fortran front-end with added F2008, Real 128 features
- Debugging support for clang/clang++ DWARF5 standards
- Debugging support for Flang DWARF4 standards
- Improved debugging support for FORTRAN arrays and pointers.

# Software Tools



## AMD PRO security

Strength at the core. Help protect your business data and personal information with AMD PRO processors.


[LEARN MORE](#)



## AMD PRO manageability

Don't get trapped inside. Cross-platform manageability means you're free to deploy your way.


[LEARN MORE](#)



## AMD PRO technologies

Built for modern business, AMD PRO technologies provide layers of protection, seamless manageability, and reliable longevity.

[LEARN MORE](#)



## AMD Infinity Guard

Help minimize potential attack surfaces as software is booted, executed, and processes your critical data



## AMD RDNA 2 Architecture


RDNA 2 architecture is the foundation for next-generation PC gaming graphics, PlayStation 5 and Xbox Series X consoles



## AMD FidelityFX

Open-source image quality toolkit for game developers


# Software Tools



## AMD Smart Access Memory

Smarter performance for great devices


[LEARN MORE](#)



## AMD CDNA Architecture

Delivers ground-breaking technologies to fuel the convergence of HPC and AI in the era of Exascale.

[LEARN MORE](#)



## AMD SmartShift Technology


Smarter Performance for Amazing Devices

[LEARN MORE](#)




## AMD FreeSync Premium Pro

AMD FreeSync™ Premium Pro raises the bar to the next level for vivid visuals, enabling an exceptional tone-mapped HDR gaming experience.



## AMD Link

Radeon™ gaming on your mobile device, anywhere




## Introduction to Semiconductors

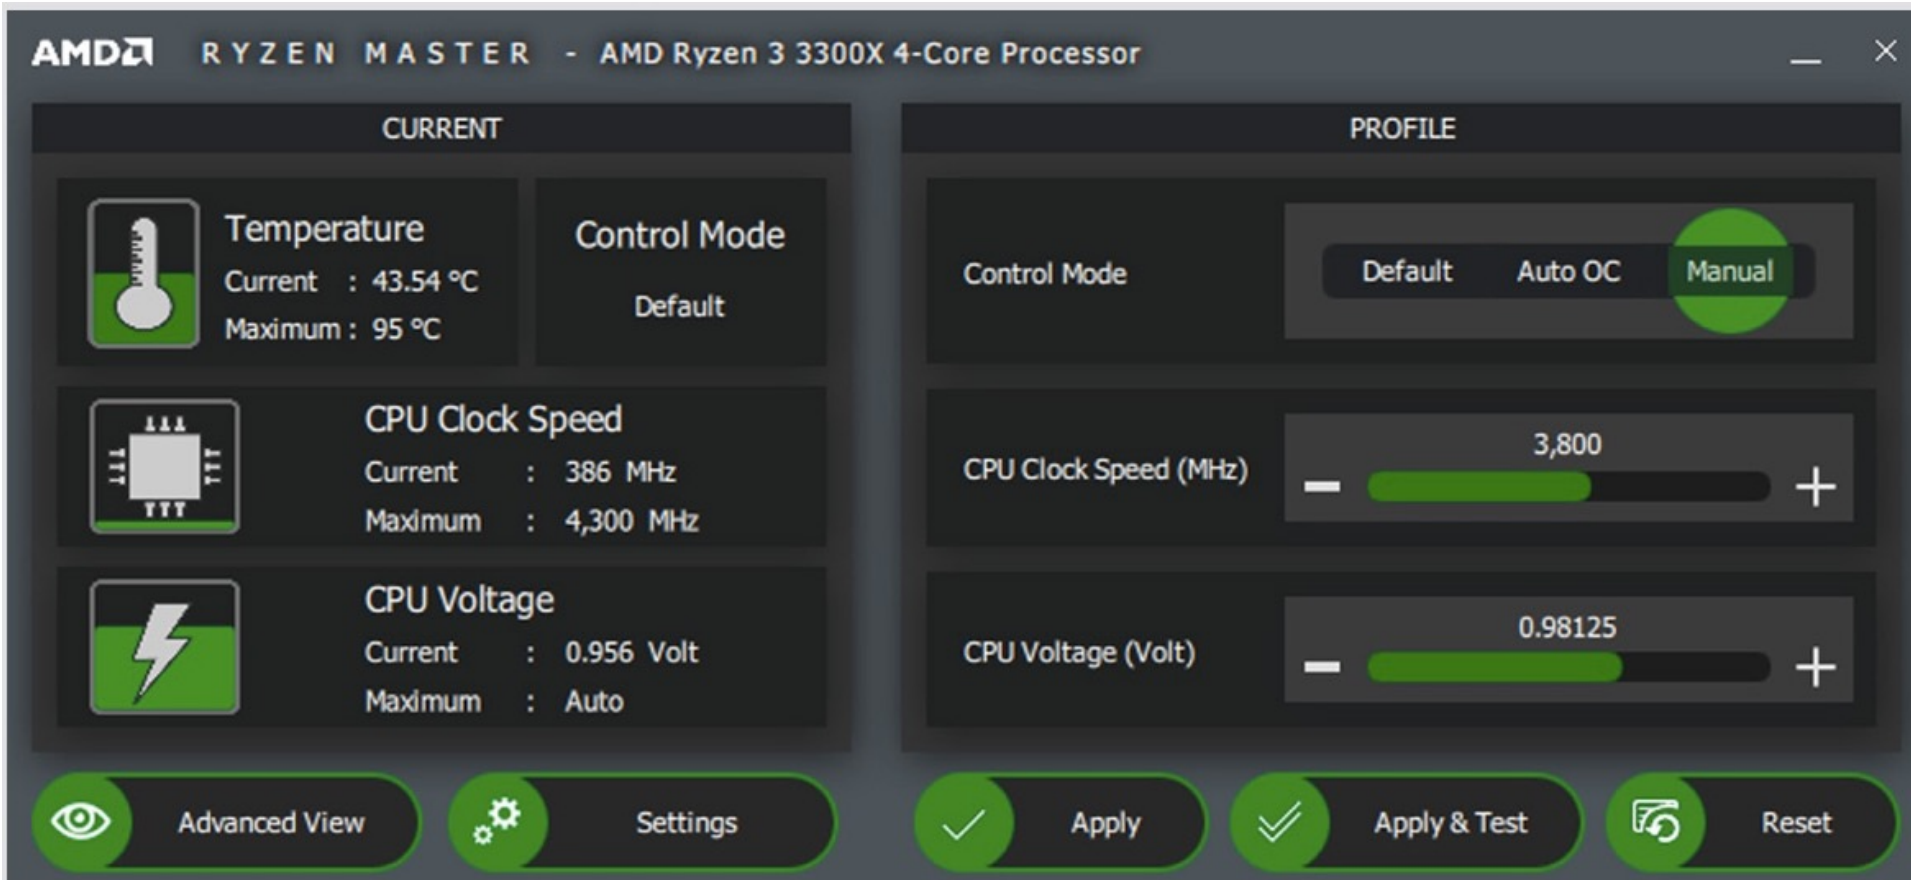
The Brains of Modern Electronics



## AMD Ryzen Master Utility for Overclocking Control



Every AMD Ryzen processor is multiplier-unlocked from the factory, so you can personalize performance to your taste. AMD provides the AMD Ryzen Master utility to access this powerful advantage.<sup>1,2</sup> As AMD Ryzen Master has evolved to support an increasingly diverse set of CPU products and features, the user interface has also grown increasingly complex. We have developed a useful “basic view” that provides access to the most essential features and telemetry. You can toggle between the new “basic view” and “advanced view” to see the overclocking features that are right for you.



## AMD Ryzen Master (Basic View)

The new basic view provides you with the ability to automatically overclock your CPU, manually control how much overclocking you wish to apply as well and give you a view of important system parameters like the CPU temperature, speed and voltages.



# Software Tools – Ryzen Master

**AMD RYZEN MASTER - AMD Ryzen 9 3900X 12-Core Processor**

Temperature, Speed, Power & Current

- Temperature: 41.88 °C (Limit 95 °C)
- Peak Speed: 1.017 GHz
- PPT (CPU): 14 % of 142 W
- TDC (CPU): 2 % of 95 A
- EDC (CPU): 63 % of 140 A

Histogram

Speed

Current Peak Avg

- Speed
  - Core 2: 924 MHz
  - Core 3: 925 MHz
  - Core 4: 924 MHz
  - Core 5: 921 MHz
  - Core 6: 923 MHz
  - Core 7: 929 MHz
  - Core 8: 931 MHz
- Temperature

Control Mode

OC Mode: Default

Cores Section

Active CCD Mode: 2

CCD	CCX	Core	Frequency	Core	Frequency	Core	Frequency
CCD 0	CCX 0	C 01	927	C 02	924	C 03	924
	CCX 1	C 04	924	C 05	921	C 06	921
CCD 1	CCX 0	C 07	929	C 08	931	C 09	921
	CCX 1	C 10	926	C 11	925	C 12	1010

Voltage Control

CPU Voltage: 1.475

Additional Control

Simultaneous Multithreading: ON

Legacy Compatibility Mode: OFF

# Software Tools – Ryzen Master



**AMD RYZEN MASTER - AMD Ryzen 5 3400G with Radeon Vega Graphics**

Control Mode: **Included**

Default Precision Boost Overdrive Auto Overclocking

Boost Override CPU	Boost Override APU GFX	PPT	TDC (CPU)
100	200	1000	60
EDC (CPU)	TDC (SOC)	EDC (SOC)	
168	60	90	

Cores Section: **Included**

CCX 0	C 01 3700	C 02 3700	C 03 3700	C 04 3700
-------	-----------	-----------	-----------	-----------

Voltage Control: **Included**

CPU Voltage: 0.96875 Volt

Additional Control: **Included**

Simultaneous Multithreading: **OFF ON**

APU GFX Speeds: **Included**

APU GFX Clock: 1450 APU GFX Voltage: 1.1

Memory Control: **Included** Push All Auto Push All Manual

Memory Clock: 1600

Voltage Controls:

MEM VDDIO: Auto	SOC Voltage: Auto	CLDO VDDP: Auto
-----------------	-------------------	-----------------

DRAM Timing Configuration:

CAS Latency: Auto	Row Precharge Delay: Auto	Read Row-Column Delay: Auto	Write Row-Column: Auto
Row Cycle Time: Auto	RAS Active Time: Auto	CAS Write Latency: Auto	Trfc: Auto
Trfc2: Auto	Trfc4: Auto	Tfaw: Auto	TrrdS: Auto
TrrdL: Auto	Twr: Auto	Twrs: Auto	TwtrL: Auto

## AMD ROCm™ Open Software Platform

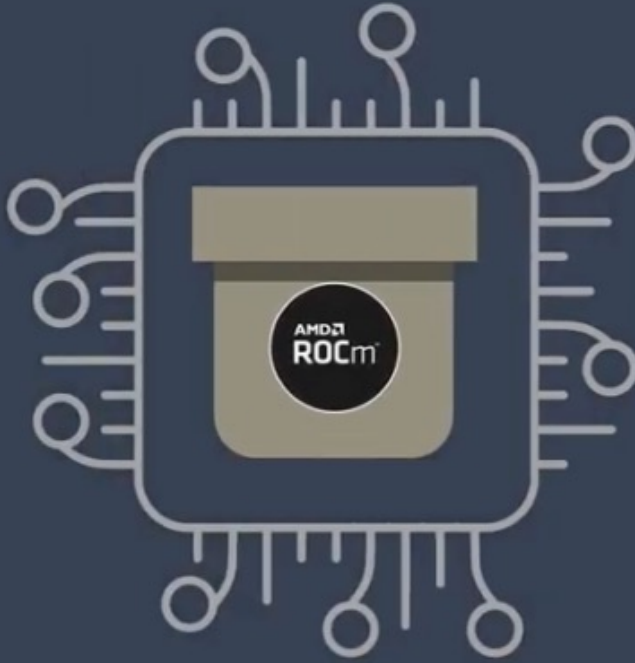
Code once, use it everywhere

The AMD CDNA architecture is built on AMD ROCm™ open software platform, the industry's only open software platform for GPU computing. The industry no longer faces a singular option that locks them into one vendor. AMD ROCm offers an open and portable platform for those looking to unleash their discoveries.





**AMD**   
**ROCm**


# ROCM

## What is ROCm™?



AMD ROCm

-  TensorFlow
-  PYTORCH
-  KOKKOS
-  RAJA

 Subscribe


# ROCM

## What is ROCm™?

**ACCELERATED COMPUTE**

**OPEN SOURCE NATIVE**

**MULTI-PLATFORM ENABLED**



**HPC & ML FRAMEWORK  
INTEGRATION**

**CONTAINER READY**

**OPEN PROGRAMMING MODEL**



**Open Source Exascale-class Platform for  
Accelerated Compute**

# Section



# News

*AMD News in AMD*

**AMD News Report WW47**

- Announcements:
  - AMD and IBM [announced](#) a joint development agreement to advance confidential computing and artificial intelligence for the cloud. The announcement was covered by multiple outlets including [Anandtech](#), [Marketwatch](#), [ZDnet](#) and others.
  - AMD [announced](#) the new Ryzen Embedded V2000 processor. Coverage has appeared in outlets including [CNX Software](#), [HotHardware](#), [Serve The Home](#), [Tom's Hardware](#) and others.
- Press continued to cover the launch of the Ryzen 5000 Series desktop processors and post review content, praising the new processors' overall performance, the architectural advancements of "Zen 3" and AMD's dominance in the CPU market. New coverage included [Gamers Nexus](#), [Hardware Canucks](#), [MMORPG](#), [SemiAccurate](#) and others, with SemiAccurate concluding, "The 19% IPC uplift means that AMD wins at everything now."
- [Tom's Hardware](#) and others reported that the Ryzen 9 5950X set an overclock world record, hitting 6.362 MHz on all 16 cores.
- A number of press, including [HEXUS](#), [PCWorld](#), [TechRadar](#) and [Tom's Hardware](#), discussed the availability of the Ryzen 5000 Series given how quickly many of the SKUs sold out. Most press argued that just because the CPUs had high demand does not mean this was a "paper launch."
- [PCMag](#) reviewed the ASUS TUF gaming A17 highlighting how the system is "a welcome change from low-cost gaming laptops with a skimpy 256GB of storage, and its Ryzen 7 processor is a real powerhouse."
- [ServeTheHome](#) reviewed the AMD EPYC 7H12 processor, naming it "the fastest CPU you can get in 2020" for raw computational power.

## AMD News in AMD

### AMD News Report WW46

- Announcements:
  - Following the reviews embargo lift of the AMD Ryzen 5000 Series desktop processors, sentiment has been overwhelmingly positive, with press praising the lineup's performance, the improvements of the "Zen 3" architecture and AMD's dominance at the top of the CPU market in gaming and content creation. Coverage appeared in [AnandTech](#), [Forbes](#), [HEXUS](#), [Linus Tech Tips](#), [PCWorld](#), [TechRadar](#), [TweakTown](#) and more.
    - [PCWorld](#): "many will see today as an historic shift in computing power...AMD has finally knocked Intel to the floor, and is raising its boxing gloves in victory as the flash bulbs pop and the ref declares a winner."
    - [ExtremeTech](#): "'Zen 3' is an unparalleled success for AMD," as the new architecture "has literally redefined what kind of performance is possible within a given desktop power envelope."
    - [JayzTwoCents](#): "the 5900X is quite honestly the best CPU you could possibly buy right now...for gaming, content creation, productivity, value – it's all there!"
  - Lawrence Livermore National Labs announced the new AMD EPYC-based "Mammoth" supercomputer system for memory intensive research workloads, including COVID-19 simulations and analysis. News appeared on [The Next Platform](#), [HPC Wire](#) and [Silicon Angle](#).
  - The reviews embargo lifted on Xbox Series X and S consoles, harnessing the power of "Zen 2" CPU cores and RDNA-based graphics to usher in the next generation of high-performance console gaming. AMD coverage sentiment has been positive, with press positioning the all-AMD hardware consoles as leaders in both raw performance and stunning, high fidelity visuals. Coverage appeared in [IGN](#), [WIRED UK](#), [PCMag](#) and more.
  - Reviews of the PlayStation 5 began to appear, with positive sentiment around AMD hardware and press praising the faster load times, better frame rates and gorgeous visuals enabled by the all-AMD console platform. Coverage appeared in [The Verge](#), [PCGamer](#), [CNN](#) and more.



## AMD

### *AMD News in AMD*

### **AMD News Report WW44**

#### **AMD news:**

- Fortune published its annual Most Powerful Women in Business list, [featuring](#) Lisa Su. Aaron Pressman [covered](#) how the most powerful women in tech have handled the pandemic, highlighting Lisa for her promotion of greater diversity and representation in her sector, as AMD “rocketed ahead this year as Intel stumbled.”
- Lenovo [announced](#) the Ryzen-powered Lenovo Legion Slim 7, which will be available starting this month.
- Ian Cutress posted an interview with Mark Papermaster diving deeper into “Zen 3” and the Ryzen 5000 Series processors on [AnandTech](#) and his YouTube channel, [TechTechPotato](#).
- [CNET](#) posted a positive review of the AMD-powered Dell G5 15 SE, highlighting how the system “delivers relatively high performance for the money.”
- [UltrabookReview](#) reviewed the Ryzen-powered Lenovo ThinkPad T14s, concluding, “AMD’s Ryzen Pro hardware is highly competent in all sorts of loads and highly efficient at the same time, allowing for excellent performance even in this slim ThinkPad T14s chassis.”
- AMD EPYC CPUs and Instinct GPUs were included in a number of newly announced HPC systems from HPE customers. These include the [550 petaflop CSC Lumi system](#), the [50 petaflop Pawsey system](#), the all [EPYC Los Alamos National Laboratory Chicoma system](#) and the [IT4 system in the Czech Republic](#). Coverage of the systems appeared in multiple outlets including [HPCwire](#), [The Next Platform](#), [VentureBeat](#) and others.
- [insideHPC](#) reported on this year’s Exascale Day, sharing an interview with Brad McCredie of AMD and other HPC industry executives. Brad highlighted the AMD EPYC and Instinct processors that will power the upcoming Frontier and El Capitan supercomputers.
- [Nextgov](#) discussed the efforts made by AMD, HPE, the Energy Department’s Exascale Computing Project or ECP, Lawrence Livermore, Argonne and Oak Ridge National Labs to share information about the potential capabilities of exascale systems.

# Nvidia News

- Nvidia news:

- [Bloomberg News](#) discussed Nvidia's acquisition of Arm and the many questions around the deal.
- Huawei and other Chinese firms [are reportedly raising concerns](#) with regulators over Nvidia's acquisition of Arm as they fear Nvidia could force Arm to stop working with them. The companies are asking for the deal to be rejected, or for regulators to insist on specific conditions.
- [VideoCardz](#) reported that Nvidia is allegedly preparing a new GeForce RTX 30 "GA102-150" Series GPU to take on the Radeon RX 6800 Series. The SKU would supposedly feature 7,424 CUDA cores, a 320-bit memory bus and 10GB of GDDR6 memory.
- [VideoCardz](#) reported that Nvidia allegedly told board partners it will not launch the RTX 3080 20GB and RTX 3070 16GB graphics cards as planned. Originally slated for a December launch, this information from two independent sources noted the SKUs were cancelled, not postponed.
- [RedGamingTech](#) claimed several sources have shared timelines for upcoming Nvidia RTX 30 Series SKUs: the RTX 3060 Ti launch is allegedly confirmed for November 17; the RTX 3060, a 6GB midrange card, should launch around January 7; and 4GB and 6GB RTX 3050 variants will launch in mid-Q1 2021 for the entry-level market.
- Nvidia [updated](#) its NVDEC (video decoding) and NVENC (encoding) matrixes for the RTX 30 Series, confirming that Nvidia GA102 and GA104 GPUs will decode the new AV1 open format up to 10-bit and 8Kx8K resolution. The cards will not support AV1 encoding and only the HEVC video format is supported.
- [VideoCardz](#) tweeted that RTX 3070 Founders Edition reviews will go live on October 27, and AIB reviews will post on October 29.
- Nvidia [announced](#) new AI-powered Neural Filters in Adobe Photoshop, including the new Smart Portrait Filter, allowing creatives to make complex adjustments to photos leveraging Nvidia RTX GPUs.

-- Contributed by AMD Communications on 27 October 2020

# AMD Buying Xilinx?

## News Article

### AMD News in Our AMD **More On AMD - Xilinx**

Article by our own Kevin Krewell

Once again, the chip company acquisition rumors are back! In this acquisition rumor AMD is in negotiations to buy Xilinx. And, again, this story comes from the Wall Street Journal, just as the rumor of NVIDIA buying Arm was, the veracity of which was later substantiated. The reason why these speculative reports come from the Wall Street Journal is that it is well connected with the investment banking community, which gets involved in these types of large acquisition deals. That said, AMD acquiring Xilinx is something that really seems to have come out of left field to us.

More: [AMD and Xilinx: A Match Made in Silicon Valley?](#)

*From EE Times -- Contributed by David Laws on 12 October 2020*

## AMD Buying Xilinx?

Advanced Micro Devices Inc. is in advanced talks to buy rival chip maker Xilinx Inc., according to people familiar with the matter, in a deal that could be valued at more than \$30 billion and mark the latest big tie-up in the rapidly consolidating semiconductor industry.

# AMD Buying Xilinx?

---



❖ AMD buys Xilinx for \$30B!

## Xilinx

From Wikipedia, the free encyclopedia

**Xilinx**, Inc. (/ˈzaɪlɪŋks/ *ZY-links*) is an American **technology company** that develops highly flexible and adaptive processing platforms. The company invented the **field-programmable gate array** (FPGA), programmable system-on-chips (SoCs), and the adaptive compute acceleration platform (ACAP). It is the **semiconductor company** that created the first **fabless manufacturing** model.<sup>[4][5]</sup> Xilinx's products are used across many industries and technologies, including the data center, wired & wireless communications, AI/ML, automotive, industrial, consumer, aerospace and defense and Broadcast & Pro-AV.

Co-founded by Ross Freeman, Bernard Vonderschmitt, and James V Barnett II in 1984, the company went public on the NASDAQ in 1989.

AMD announced its acquisition of Xilinx in October 2020.<sup>[6]</sup>

# Xilinx

## Xilinx, Inc.



Xilinx headquarters in the United States

<b>Type</b>	Public
<b>Traded as</b>	NASDAQ: XLNX <a href="#">↗</a> NASDAQ-100 Component S&P 500 Component
<b>Industry</b>	Integrated circuits
<b>Founded</b>	1984; 36 years ago <sup>[1]</sup>
<b>Founder</b>	Jim Barnett Ross Freeman Bernie Vonderschmitt
<b>Headquarters</b>	San Jose, California, U.S.
<b>Area served</b>	Worldwide

### Key people

- Dennis Segers (chairman of the board)
- Victor Peng (president, CEO)
- Brice Hill (CFO)<sup>[2]</sup>
- Ivo Bolsens (senior vice president, CTO)
- Kevin Cooney (senior vice president, CIO)
- Catia Hagopian (senior vice president, general counsel)
- Vincent L. Tong (executive vice president, global operations and quality)
- Liam Madden (executive vice president, hardware and systems product development)
- Matt Poirier (senior vice president, corporate development and investor relations)
- Salil Rajee (executive vice president, software and IP products)
- Marilyn Stiborek Meyer (senior vice president, global human resources)
- Mark Wadlington (senior vice president, global sales)

### Products

FPGAs, CPLDs

### Revenue

- ▲ US\$ 3.16 billion (2020)<sup>[3]</sup>
- ▲ US\$ 3.06 billion (2019)<sup>[3]</sup>

### Operating income

- ▼ US\$ 791.888 million (2020)<sup>[3]</sup>
- ▲ US\$ 956.799 million (2019)<sup>[3]</sup>

### Net income

- ▼ US\$ 792.721 million (2020)<sup>[3]</sup>
- ▲ US\$ 889.750 million (2019)<sup>[3]</sup>

### Total assets

- ▼ US\$ 4.693 billion (2020)<sup>[3]</sup>
- ▲ US\$ 5.151 billion (2019)<sup>[3]</sup>



## Our AMD

*AMD News in Our AMD*

### Rumor Mill - New AMD Graphic Card

New rumors have emerged that suggest AMD is working on a seriously powerful 'Big Navi' graphics card that will take on Nvidia's new RTX 3080, as well as the GPU found in the upcoming Xbox Series X console.

What's most noticeable about those alleged specs is the 2.4GHz clock

More: [AMD Radeon RX 6900 XT Big Navi GPU could crush the Nvidia RTX 3080 and Xbox Series X](#)

## Science and Industry

*Semiconductors in Science and Industry*

### Intel's Fading Crown

Intel Corp. is still the dominant provider of chips for personal computers and servers, but its reign as the chip king of Silicon Valley has ended, and the unresolved question is whether its woes stem from reaching the limits of technology, the people in charge or a mixture of both.

More: [How did Intel lose its Silicon Valley crown?](#)

## *AMD News in Our AMD*

### **AMD News Report WW41**

- Announcements:
  - AMD [announced](#) VMware vSphere now supports the advanced SEV-ES security feature. AMD also highlighted the growth of the EPYC-based VMware ecosystem. Coverage appeared in [CRN](#), [SDXcentral](#), [Storage Review](#), [ZDNet](#) and others.
  - AMD [announced](#) the Okinawa Institute of Science and Technology Graduate University is using AMD EPYC 7702 processors to power a 2.36 petaflop system used for research at the university.
  - AMD announced the new [Radeon Software Adrenalin 2020 Edition 20.9.2](#) will deliver Day-0 support for STAR WARS: Squadron, bringing impressive performance and high-framerate gameplay from 1080p to 4K resolutions with [Radeon RX 5000 Series](#) graphics.
  - AMD announced the [AMD Radeon Graphics 'Raise the Game' offer](#) has been extended through December 5, 2020, giving Radeon gamers more time to score some of 2020's hottest titles with the purchase of select AMD Radeon graphics products, including Godfall (Radeon RX 5500 series ) or both Godfall and World of Warcraft: Shadowlands (Radeon RX 5600 and 5700 cards).