



Intel Processors Over the Years

Intel processors have been on the cutting edge of technology for over half a century, underpinning the creation of today's digital global economy.



Written by: **Neil Cumins**, Senior Analyst
Updated Aug 08, 2024

Editor Reviewed: **Sandra Mardenfeld**, Senior Editor

Business News Daily earns compensation from some listed companies. [Editorial Guidelines](#).

Intel Going Forward

Over the last 56 years, Intel Corporation has played a central role in the computing sector. Founded in California long before the state became the spiritual home of tech firms, Intel is now the world's largest semiconductor chip manufacturer. However, the big numbers surrounding this global tech giant (over 131,000 employees, \$146.7 billion net worth) are underpinned by tiny products — semiconductor chips that serve as computer processors.

Without processors, computers wouldn't work. Intel has been the dominant force in developing the global computing industry, the growth of the internet and modern-day reliance on cloud services. But while Intel's story is well known, the history of its processors is less extensively documented.

Did You Know?



According to [Statista](#), Intel and AMD hold 96 percent of the global computer central processing unit (CPU) market.

History of Intel Processors

To celebrate the development of products that have literally changed the world, here's a walkthrough of the history of Intel processors, starting with the first commercially available processor.

1971–81: The 4004, 8008 and 8800

The 4004 was the first complete CPU on a single chip, packaged in a 16-pin ceramic dual in-line package. The 4004 was released initially with a clock speed of 108 kHz (and scaled up to 740 kHz). Produced in a 10 µm (10,000 nm) process, the 4004 had 2,300 transistors and delivered a performance of 0.07 MIPS.

The 8-bit 8008 replaced the 4004 in 1972 with a 0.5 MHz to 0.8 MHz clock speed and 3,500 transistors. It was primarily used in the TI 742 computer. The 8080 followed in 1974 with 4,500 transistors in 6,000 nm and up to 2 MHz. It became famous for being used in the Altair 8800 and Boeing's AGM-86 cruise missile.

None of these chips sold in considerable volumes.

Tip



If you're already feeling overwhelmed with technical jargon, check out our [guide to key technology terms](#) for explanations of some terms in this article.

1978–82: iAPX 86 (8086), 8088 and 80186 (16-bit)

The 8086, also known as the iAPX 86, was Intel's first commercial 16-bit CPU and is considered the chip that launched the era of x86 processors. With 29,000 transistors built in a 3,000 nm design, the 8086 was clocked from 5 MHz to 10 MHz and achieved up to 0.75 MIPS in computers, such as the IBM PS/2.

The IBM 5150, the first PC, came with the 8088 (5-8MHz), which was identical to the 8086 except for its 8-bit internal bus. In 1982, Intel launched the 80186 CPU, which was also based on the 8086 but was built in 2,000 nm and hit more than 1 MIPS at a 6 MHz clock speed. The Tandy 2000 was among the first PCs that used the 80186.

1981: iAPX 432

The iAPX 432 is one of the few Intel processor designs that flopped and Intel does not talk about it anymore. Other future ill-fated processor designs include the i860/i960 of the early 1990s and the highly integrated Timna processor of 2000.

Introduced in 1981, the 432 was Intel's first 32-bit design — an amazingly complex design for its time that integrated hardware-based multitasking and memory management features.

Designed for high-end systems, the downfall of the 4-8 MHz 432 was that it was much more expensive to produce and slower than the emerging 80286 design.

While the 432 was initially designed to replace the 8086 series, the project ended in 1982.

1982: 80286

Intel's 80286 debuted with memory management and wide protection abilities. In 1991, it reached clock speeds up to 25 MHz and performed more than 4 MIPS. This processor was popular in IBM-PC AT and AT PC clones. The chip was manufactured at 1,500 nm and included 134,000 transistors.

The 80286 is remembered as the Intel processor that provided the highest performance gain over its predecessor and one of the most cost-efficient processors Intel ever produced. In 2007, Intel stressed that only the new Atom processor was about as cost-efficient as the 80286 25 years earlier.

1985–94: 386 and 376

The 32-bit era began with the release of the 386DX CPU in 1985. With 275,000 transistors (1,500 nm) and clock speeds ranging from 16 MHz to 33 MHz, the CPU hit up to 11.4 MIPS.

In 1988, Intel followed up with the 1,000 nm 386SX, which had a narrower 16-bit bus to target mobile and low-cost desktop computing systems. Although the 386SX remained fully 32-bit capable internally, the data bus was cut to 16 bits to simplify the circuit board layout and reduce costs. Additionally, although not critical at the time, only 24 pins were connected to the 386SX's address bus, which effectively limited it to addressing 16MB of memory.

Both chips lacked a math coprocessor. Due to early problems with the i387 coprocessor not being production-ready in time for the 80386, both chips had to use the 80287 as their math coprocessor until the 80387 was released to the market.

Intel's first notebook chip, the 386SL, arrived in 1990 as a highly integrated design with an on-chip cache, bus and memory controller. The processor had 855,000 transistors and ran between 20 MHz and 25 MHz. The 376 (1989) and 386EX (1994), both for embedded systems, completed the 376/386 processor family.

Despite becoming obsolete as a personal computer CPU in the early 1990s, Intel continued to manufacture the 80386 family until September 2007 due to market demand for the chip in embedded systems and its wide use by the aerospace industry.

1989: 486 and i860

The 486, designed under the guidance of Pat Gelsinger, former CEO of VMware, drove Intel through its greatest growth phase. The 1,000 nm and 800 nm designs were launched as the 486DX with 25 MHz to 50 MHz, which included 1.2 million transistors and delivered 41 MIPS. The low-end 486SX (a 486DX with a disabled math coprocessor) followed in 1991 with 16 MHz to 33 MHz.

In 1992, Intel introduced an update as the 486DX2 (SX2) with up to 66 MHz while the 486SL as an enhanced 486SX was offered for notebooks (up to 33 MHz, 800 nm, 1.4 million transistors). The final stage of the 486 series was the 486DX4 with up to 100 MHz, which was marketed as an economical solution for those who did not want to spend more money on the new Pentium systems. The DX4 was built in a 600 nm process, had 1.6 million transistors and was rated at 70.7 MIPS.

The year 1989 was also the release year of the i860, Intel's attempt to enter the RISC processor race and the company's second major shot at the high-end computer segment. The i860 and i960 never succeeded and were canceled in the early 1990s.

Tip



For those entering the information technology (IT) field, the **best IT certifications** include the **CompTIA A+ certification**, the **ACMT (Apple) certification** and the **BICSI Technician certification**.

1993: Pentium (P5, i586)

The original Pentium was introduced in 1993. In 2005, there were rumors that Intel would drop the name in favor of the new Core brand, but the Pentium brand lives on. The brand is an essential part of Intel's history and a departure from the 286/386/486 processor numbers. Intel reportedly chose a word to be able to protect the trademark against AMD, which also offered 486-labeled processors.

The P5 Pentium launched with 60 MHz in 1993 and was available with up to 200 MHz (P54CS) in 1996. The original 800 nm design had 3.1 million transistors but scaled to 3.3 million in the 350 nm 1996 design. The P55C was announced in 1997 with multimedia extensions (MMX), expanding the processor design to 4.5 million transistors and a 233 MHz clock speed. The mobile version of the Pentium MMX remained available until 1999 and reached 300 MHz.

1994–99: Bumps in the road

Throughout the years, Intel has released many successful additions to its lineup of processors and architectures but not without running into the occasional bump in the road.

In 1994, a professor at Lynchburg College discovered a bug in the Intel P5 Pentium floating-point unit that affected several models of the original Pentium processor. The bug, known as the Pentium FDIV bug, caused the processor to return incorrect decimal results in certain division operations, which created issues in fields like mathematics and engineering, where precise results were needed.

Although rare, *Byte* magazine estimated that about one in 9 billion divides would produce incorrect results. Intel attributed the flaw to missing entries in the processor's lookup table used by floating-point division circuitry.

In 1999, Intel released the Pentium III processor, the first x86 processor to feature a unique ID number dubbed the PSN, or processor serial number. The PSN could be readily accessed by software if not disabled by the user in the BIOS through the use of the CPUID instruction.

After its discovery, the PSN caused Intel to come under fire from a number of groups, including the European Parliament, which cited privacy concerns over the ability of PSN to be used by surveillance groups to identify individuals. Intel subsequently removed the PSN feature from its future processors, including the Tualatin-based Pentium IIIs.

1995: Pentium Pro (P6, i686)

Upon its release, the Pentium Pro was a largely misunderstood processor. Many believed the Pro was intended to replace the P5. However, as a precursor to the Pentium II Xeon, the Pentium Pro was tailored to deal with workloads typical for servers and workstations.

Other than what the name implies, the Pentium Pro's architecture was different from the regular Pentiums and supported out-of-order execution, for example. In addition to the different architecture, the Pentium Pro had a 36-bit address bus, which supported up to 64GB of memory.

The Pentium Pro was built in 350 nm, had 5.5 million transistors and came in several variants with clock speeds ranging from 150 MHz to 200 MHz. Its most famous application was its integration into the ASCI Red supercomputer, which was the first to break through the 1 teraflop performance barrier.

Key Takeaway



The Intel Xeon range was designed for nonconsumer products like servers and workstations. As such, Xeon products tend to have higher core counts, greater cache memory and extra reliability, availability and serviceability features for stable operation.

1997: Pentium II and Pentium II Xeon

The Pentium II was a consumer-focused processor developed on top of the sixth-generation P6 architecture. It was the first Intel CPU delivered in a cartridge-like slot module instead of a socket device. The Pentium II had 2 million more transistors (7.5 million) than the P6, significantly improving 16-bit execution, which was a problem in the initial P6 release and carried on the MMX instruction set that was introduced with the Pentium.

The Pentium II was released with the 350 nm Klamath core (233 MHz and 266 MHz). Deschutes arrived in 1998 as a shrink to 250 nm and clock speeds up to 450 MHz. It was also offered as Pentium II Overdrive as an upgrade option for the Pentium Pro. Mobile Pentium II processors got the 250 nm Tonga, 250 nm and 250 nm/180 nm Dixon cores.

In the same year, Intel also offered the Deschutes core as a Pentium II Xeon with a larger cache and dual-processor support.

1998: Celeron

While Celerons are based on the company's current processor technology, they usually come with substantial downgrades, such as less cache memory, which positions them as processors that are "good enough" for the most basic PC applications. Their presence allows Intel to compete at the bottom end of the PC market.

The first Celeron series was based on the 250 nm Covington core for desktops and the 250 nm Mendocino core (19 million transistors, including L2 on-die cache) for notebooks. The processors were available from 266 MHz to 300 MHz on the desktop and up to 500 MHz on the mobile side. They were updated well into the days of the succeeding Pentium III. Today's Celerons are based on Sandy Bridge architecture.

Key Takeaway



Intel's low-end consumer processor, Celeron, launched in 1998 as a variant of the Pentium II processor and remains popular over 25 years later.

1999: Pentium III and Pentium III Xeon

Released in 1999, the Pentium III was Intel's initial contender in the gigahertz race with AMD. The CPU also entered Transmeta's low-power challenge in early 2000. The chip was initially released with the

250 nm Katmai core and was quickly scaled down to 180 nm with Coppermine and Coppermine T and 130 nm with the Tualatin core.

Due to the integrated L2 cache, the transistor count jumped from 9.5 million in Katmai to 28.1 million in the following cores. The initial clock speed was 450 MHz and eventually reached 1,400 MHz with Tualatin. Intel was criticized for rushing out the first gigahertz versions to compete with AMD's Athlon, which forced the company to recall its gigahertz processors and re-release them later.

Also noteworthy on the consumer side was the announcement of the Mobile Pentium III in 2000, which introduced SpeedStep and a processor clock speed scaling ability, depending on its operation mode. The Mobile Pentium III was announced one day before the announcement of the Transmeta Crusoe processor and many still believe that the Mobile Pentium III would not have been released without the pressure of Transmeta, which was famous for employing Linux inventor Linus Torvalds.

The Pentium III Xeon was the last Xeon processor tied to the Pentium brand. The chip was released with the Tanner core in 1999. Controversially, Intel introduced the PSN with the Pentium III. The feature caused several privacy complaints and Intel eventually removed the feature and did not carry it over to future CPUs.

2000: Pentium 4

The Pentium 4 arguably took Intel on a path that led to the most dramatic transformation in the company's history. Launched in 2000 with the 180 nm Willamette core (42 million transistors), the chip's Netburst architecture was designed to scale with clock speed; Intel envisioned that the foundation would allow the company to hit frequencies of more than 20 GHz by 2010. Netburst, however, was more limited than initially thought and, by 2003, Intel knew the current leakage and power consumption was increasing too rapidly with higher clock speeds.

Netburst launched with 1.3 GHz and 1.4 GHz, increased to 2.2 GHz with the 130 nm Northwood core (55 million transistors) in 2002 and to 3.8 GHz with the 90 nm Prescott core (125 million transistors) in 2005. Intel also launched the first Extreme Edition processors with the Gallatin core in 2003.

Over time, the Pentium 4 series became increasingly confusing, with Mobile Pentium 4-M processors, Pentium 4E HT (hyperthreading) processors with support for a virtual second core and Pentium 4F processors with the 65 nm Cedar Mill core (Pentium 4 600 series) in 2005.

Intel planned to replace the Pentium 4 family with the Tejas processor but canceled the project when it was clear that Netburst would not be able to reach clock speeds beyond 3.8 GHz. Core, the following architecture, was a dramatic turnaround to much more efficient CPUs with a strict power ceiling that put Intel's gigahertz machine in reverse.

2001: Xeon

The first Xeon that did not bring the Pentium brand along was based on Pentium 4's Netburst architecture and debuted with the 180 nm Foster core. It was available with 1.4 GHz to 2 GHz clock speeds.

The Netburst architecture continued until 2006, when Intel expanded Xeon to a full line of UP and MP processors with the 90 nm Nocona, Irwindale, Cranford, Potomac and Paxville cores, as well as the 65 nm Dempsey and Tulsa cores.

Similar to its desktop processors, the Netburst processors suffered from excessive power consumption, which forced Intel to revise its processor architecture and strategy. The Netburst Xeons died with the dual-core Dempsey CPU, which had a clock speed of up to 3.73 GHz and 376 million transistors.

Today's Xeons are still based on the technology foundation that is also used for desktop and mobile processors, but Intel keeps them in a tight power envelope. The 2006 dual-core Woodcrest chip, a variant of the desktop Conroe chip, was the first representative of this new idea.

The current Xeons are based on 32 nm Sandy Bridge and Sandy Bridge-EP architecture and Westmere processor designs. The CPUs have up to 10 cores, clock speeds up to 3.46 GHz and up to 2.6 billion transistors.

2001: Itanium

The Itanium has been Intel's most misunderstood processor, yet it survived over a long period. While it follows the idea of the i860 and iAPX 432, it has found some powerful supporters and hasn't been cut yet. The processor was launched as Intel's first 64-bit processor and was believed to be Intel's general idea for a 64-bit platform. However, the Itanium suffered in the 32-bit department and was heavily criticized for its lack of performance in this segment.

Itanium was launched in 2001 with the 180 nm Merced core as a mainframe processor with 733 MHz and 800 MHz clock speeds and 320 million transistors — more than six times the count of a desktop Pentium at the time.

The Itanium 2 followed in 2002 (180 nm McKinley core as well as 130 nm Madison, Deerfield, Hondo, Fanwood and Madison cores) and wasn't updated until 2010 when Intel launched the Itanium 9000 with the 90 nm Montecito and Montvale cores as well as the 65 nm Tukwila core with a massive 24MB on-die cache and more than 2 billion transistors.

2002: Hyper-Threading

In 2002, Intel released the first modern desktop processor with simultaneous multithreading technology (SMT), known as Intel Hyper-Threading (HT) Technology. HT Technology first appeared in Intel's Prestonia-based Xeon processors and later in the Northwood-based Pentium 4 processors. The operating system can execute two threads simultaneously by allowing one thread to run while the other is stalled, usually due to a data dependency.

At the time, Intel claimed a performance improvement of up to 30 percent over a nonhyperthreaded Pentium 4. Previous tests showed that a hyperthreaded 3 GHz chip could surpass the speed of a non-hyperthreaded 3.6 GHz chip under certain conditions. Intel has continued to include hyperthreading in various processors, including the Itanium, Pentium D, Atom and Core i-Series CPUs.

Key Takeaway



Hyperthreading works by duplicating certain processor sections, allowing the operating system to address a single physical processor with two logical processors per core.

2003: Pentium M

The Pentium M 700 series, launched with the 130 nm Baniyas core in 2003, was targeted at mobile computers. It embodied the philosophy of an Intel brand that no longer focused its processors on clock speed but rather on power efficiency. The processor was developed by Intel's design team in Israel, led by Mooly Eden, who held a key executive role at the firm for many years.

Baniyas dropped its clock speeds to between 900 MHz and 1.7 GHz, down from the Pentium 4 Mobile's 2.6 GHz. However, the processor was rated at just 24.5 watts TDP, while the Pentium 4 chip was at 88 watts. The 90 nm shrink was called Dothan and dropped its thermal design power to 21 watts. Dothan had 140 million transistors and clock speeds of up to 2.13 GHz.

The direct successor of Dothan was Yonah, which was released in 2006 as Core Duo and Core Solo but was not related to the Intel Core microarchitecture. The Baniyas core and its impact on Intel is seen in the same manner as the 4004, 8086 and 386.

2005: Pentium D

The Pentium D was Intel's first dual-core processor. Still based on Netburst, the first version had the 90 nm Smithfield core (two Northwood cores) and was released as the Pentium D 800 series. It was succeeded by the 65 nm Presler (with two Cedar Mill cores) dual core.

Intel also released Extreme Editions of both processors, capping the maximum clock speed at 3.73 GHz and setting a power consumption of 130 watts — the highest ever for any Intel consumer desktop processor (some server processors went up to 170 watts). Smithfield had 230 million transistors; Presler had 376 million.

2005–09: Terascale Computing Research Program

Intel's Tera-Scale Computing Research (TSCR) program started around 2005 to address the various challenges faced in scaling chips beyond four cores and experiment with improving communication within the processors themselves. The TSCR program has yielded several notable devices, including the Teraflops Research Chip and the Single-Chip Cloud Computer (SCC), both of which became significant contributors to Intel's Xeon Phi line of coprocessors.

The Teraflops Research Chip, codenamed Polaris, is an 80-core processor developed through the TSCR program. The chip features dual floating-point engines, sleeping-core technology and 3D memory stacking, among other things. Its purpose was to experiment on how to effectively scale beyond four cores on a single die and to build a chip capable of producing a teraflop of computing performance.

The SCC was a 48-core processor developed through the TSCR program. The idea behind the chip was to have several sets of separate cores that could communicate directly with each other, similar to how servers in a data center communicate.

The chip contains 48 Pentium cores in a 4 x 6, two-dimensional mesh of 24 tiles sharing two cores and 16KB of cache each. The tiles allow the cores to communicate with each other instead of sending and retrieving data from the main memory, significantly improving performance.

2006: Core 2 Duo

Core 2 Duo was Intel's response to AMD's Athlon X2 and Opteron processors, which were highly successful at the time. The Core microarchitecture was launched with the 65 nm Conroe (Core 2 Duo E-6000 series) on the desktop, Merom on the mobile side (Core 2 Duo T7000 series) and Woodcrest in the server market (Xeon 5100 series). Intel quickly followed with quad-core versions (Kentsfield Core 2 Quad series for the desktop, Clovertown Xeon 5300 series for servers).

The Core microarchitecture was preceded by one of the most significant restructurings at Intel, as well as a substantial repositioning of the company. While Conroe was developed, Intel positioned its remaining Pentium and Pentium D processors to drive AMD into an unprecedented price war in 2005 and 2006, while the Core 2 Duo processor regained the performance lead over AMD in 2006. Conroe was launched with 1.2 GHz to 3 GHz clock speeds and as a chip with 291 million transistors. The CPUs were updated with a 45 nm Penryn shrink in 2008 (Yorkfield for quad cores).

While Intel always attempted to deliver a die shrink every two years, the arrival of Core 2 Duo also marked the introduction of the company's tick-tock cadence, which dictates a shrink in uneven years and a new architecture in even years.

2007: Intel vPro

Around 2007, Intel introduced its vPro technology, which isn't much more than a marketing term for a suite of hardware-based technologies included on select Intel processors produced since then.

vPro, which is often confused with Intel's Active Management Technology (AMT), is mainly targeted at the enterprise market. It encompasses Intel technologies such as Hyper-Threading, AMT, Turbo Boost 2.0 and VT-x in a single package. For a computer to utilize vPro technology, it must have a vPro-enabled processor, a vPro-enabled chipset and a basic input/output system (BIOS) that supports vPro technology.

These are some of the major technologies vPro includes:

Intel Active Management Technology (AMT) is a set of hardware features that allows system administrators to remotely access and manage a computer even when the computer is

powered off. AMT's remote configuration technology allows basic configuration to be performed on systems that do not yet have an operating system or other management tools installed.

- **Intel Trusted Execution Technology (TXT)** verifies a computer's authenticity using the Trusted Platform Module (TPM). TXT then builds a chain of trust using various measurements from the TPM, which are used to make trust-based decisions about what software can run. This allows systems administrators to ensure sensitive data is only processed on a trusted platform.
- **Intel Virtualization Technology (VT)** is a hardware-based [virtualization](#) technology that allows multiple workloads to share a common set of resources in full isolation. Additionally, VT removes some of the performance overhead incurred by solely using software virtualization.

2008: Core i-Series

Intel's Core i3, i5 and i7 processors launched with the Nehalem microarchitecture and the company's 45 nm production process in 2008. The architecture was scaled to 32 nm (Westmere) in 2010 and provided the foundation for Intel processors covering the Celeron, Pentium Core and Xeon brands. Westmere scaled to up to eight cores, up to 3.33 GHz clock speed and up to 2.3 billion transistors.

Did You Know?



Westmere was effectively replaced by the 32 nm Sandy Bridge architecture in 2011, which shrank in 2012 to 22 nm in the Ivy Bridge generation (1.4 billion transistors for quad-core processors).

2008: Atom

Atom was launched in 2008 as a processor designed to power mobile internet devices and nettops. The initial 45 nm single chip was sold in a package with a chipset and a thermal design power as low as 0.65 watts. As netbooks became popular in 2008, the less power-efficient Diamondville (N200 and N300 series) core sold in far greater units than the Silverthorne core (Z500 series), which Intel envisioned to be its contender for the ultramobile market.

The initial Atom lacked integration and did not succeed in markets other than netbooks. Even the updated Lincroft (released in 2010 as Z600) could not change that scenario. The current Atom generation for desktop and netbook applications is the 32 nm Cedarview generation (D2000 and N2000 series, released in 2011). Intel attempted to expand Atom into other application areas, such as TVs but failed mainly due to its lack of integration.

The Atom SoC, with the Medfield core, was released in 2012. The Z2000 series is Intel's first offering for devices such as phones and tablets since its ARMv5-based Xscale core, which the company offered between 2002 and 2005.

Tip



If your business is **switching from a PC to a Mac**, ensure your crucial business software solutions and apps are available in Mac form.

2010: HD Graphics

In 2010, Intel introduced its Westmere architecture featuring on-die graphics, known as Intel HD Graphics. Previously, any computer not utilizing a discrete graphics card made use of the Intel Integrated Graphics residing on the motherboard's Northbridge chip.

With Intel's continued move from its Hub Architecture design to the new Platform Controller Hub (PCH) design, the Northbridge chip was eliminated entirely and the integrated graphics hardware was moved to the same die as the CPU. Unlike the previous integrated graphics solution, which had a poor reputation for lacking performance and features, Intel's HD Graphics once again made integrated graphics competitive with discrete graphics manufacturers through major performance increases and low power consumption.

Intel HD Graphics came to dominate the low-to-midrange device market, picking up an even more substantial share in the mobile device sector. The Intel HD Graphics 5000 (GT3) has a TDP of 15 watts, 40 execution units and a performance output of up to 704 GFLOPS.

In 2013, Intel launched its Iris Graphics and Iris Pro Graphics on a limited set of its Haswell processors as a high-performance version of HD Graphics. The Iris Graphics 5100 is largely the same as the HD Graphics 5000 but features an increased TDP of 28 watts, an increased maximum frequency of 1.3 GHz and a small increase in performance of up to 832 GFLOPS.

The Iris Pro Graphics 5200, referenced as Crystalwell by Intel, is the first of Intel's integrated solutions to have its own embedded DRAM, featuring a 128MB cache for performance improvements in bandwidth-limited tasks. In late 2013, Intel announced that the Broadwell-K series of processors would feature Iris Pro Graphics in place of HD Graphics.

2010: Many Integrated Core Architecture and Xeon Phi

Initial work on Intel's Many Integrated Core (MIC) Architecture began around 2010, drawing on technology from several earlier projects, such as the Larrabee microarchitecture, the SCC project and the Teraflops Research Chip. Intel's various MIC Architecture products, which would later come to be known as Xeon Phi, are coprocessors, which are specialized processors designed to increase computing performance by offloading processor-intensive tasks from the CPU.

In 2010, Intel debuted its first MIC Architecture prototype board, codenamed Knights Ferry. This PCIe card sported 32 cores at 1.2 GHz and four threads per core. The development board also featured

2GB of GDDR5 memory, 8MB of L2 cache, a power consumption of around 300 watts and performance exceeding 750 GFLOPS.

In 2011, Intel announced an improvement to its MIC Architecture, codenamed Knights Corner. It was made using the 22 nm process with Intel's Tri-Gate transistor technology and had over 50 cores per chip. Knights Corner was Intel's first commercial MIC Architecture product and was quickly adopted by many companies in the supercomputer industry, including SGI, Texas Instruments and Cray. Intel officially rebranded Knights Corner as Xeon Phi in 2012 at the Hamburg International Supercomputing Conference.

Intel revealed its second-generation MIC Architecture, dubbed Knights Landing, in June 2013. Intel announced that the Knights Landing products would be built with up to 72 Airmont cores with four threads per core using the 14 nm process. Additionally, Intel stated that each card would support up to 384GB of DDR4 RAM, include 8GB to 16GB of 3D MCDRAM and have TDPs ranging from 160 to 215 watts.

Xeon Phi products include the Xeon Phi 3100, Xeon Phi 5110P and the Xeon Phi 7120P, which are all based on the 22 nm process. The Xeon Phi 3100 is capable of more than 1 teraflop of double-precision floating-point performance, with memory bandwidth of 320 Gbps and a recommended price tag of less than \$2,000. At the high end of the spectrum, the Xeon Phi 7120P is capable of more than 1.2 teraflops of double-precision floating-point performance, 352 Gbps memory bandwidth and a price tag north of \$4,100.

2012: Intel SoCs

Intel's venture into the system-on-a-chip (SoC) market began around mid-2012 when the company launched its line of Atom SoCs. The earliest of these were merely lower-power adaptations of earlier Atom processors, which didn't see much success against ARM-based SoCs. Intel SoCs began to take off in late 2013 with the release of the Baytrail Atom SoCs based on the 22 nm Silvermont architecture.

Like the newly released Avoton chips for servers, the Baytrail chips are true SoCs, with all the components necessary for tablets and laptop computers. They feature TDPs as low as 4 watts. In addition to the Atom-based SoCs, around early 2014, Intel began a serious push to bring its more popular desktop architectures into the high-end tablet market by introducing the Haswell architecture Y stock keeping unit (SKU) suffix ultralow-power processors with TDPs around 10 watts.

In late 2014, Intel started releasing chips based on the Broadwell architecture, further extending its venture into the SoC market with quad-core chips featuring TDPs as low as 3.5 watts and support for up to 8GB of LPDDR3-1600 RAM.

Tip



When buying a secure business laptop, look for features like biometric security, smart

card readers and encryption.

2013: Core-i Series – Haswell

Intel updated its Core-i series of processors in 2013 with the debut of the 22 nm Haswell microarchitecture, which replaced the 2011 Sandy Bridge architecture.

With the introduction of Haswell, Intel also introduced the Y SKU suffix for its new low-power processors designed for ultrabooks and high-end tablets (10- to 15-watt TDP). Haswell scaled up to 18 cores with the Haswell-EP line of Xeon processors, which featured up to 5.69 billion transistors and clock speeds of up to 4.4 GHz.

In 2014, Intel released a refresh of the Haswell lineup called Devil's Canyon, featuring a modest boost in clock speeds and an improved thermal interface material to alleviate heat issues faced by enthusiasts and overclockers. The Broadwell die shrink in 2014 scaled down the architecture to 14 nm but did not replace the full line of Haswell CPUs, instead forgoing the inclusion of low-end desktop CPUs.

Key Takeaway



The Core-i series represented the point at which Intel began releasing generations of microprocessors, as opposed to separate models like the Pentium II, III and 4.

2015: Broadwell

With its fourth generation of modern processors, 2015 was the year when 14 nm architecture became the default. After a period of downsizing from 45 nm in 2010 to 22 nm with Haswell, Broadwell was 37 percent smaller than its immediate predecessor. Battery life could also be expanded by 1.5 hours, with faster wake times.

Other benefits of Broadwell included improved graphics performance with two-channel DDR3L-1333/1600 RAM via 1150 LGA sockets.

2015: Skylake

Sir to how Android used to have dessert-themed brands, each generation of Intel processor released sll J15 has had a lake-themed title. Skylake was the first, launched just seven months after

Broadwell but returning a 10 percent improvement in instructions per clock (IPC) thanks to microarchitecture improvements.

These chips were considerably more expensive, limiting their appeal. Their cache was slightly smaller than Broadwell's, even though speeds could reach 4 GHz. They were used exclusively in Xeon processors, whereas Broadwell had been used in Celeron, Pentium, Xeon and Core M chips.

2016: Kaby Lake

The first Intel microprocessor to abandon the company's iconic "tick-tock" manufacturing and design model, Kaby Lake was also significant for being the first Intel hardware incompatible with Windows 8 or older iterations.

Improvements over Skylake included faster CPU clock speeds and clock speed changes, although IPC figures were unchanged. Kaby Lake offered superior 4K video processing and was used in Core, Pentium and Celeron processors — but, significantly, not Xeon. A later refresh of Kaby Lake in early 2017 introduced R models supporting DDR4-2666 RAM.

2019: Ice Lake

The Ice Lake microarchitecture was introduced in 2019 as part of Intel's 10th-generation Core processors, utilizing a 10 nm process technology. This architecture brought several advancements, including support for Wi-Fi 6 and Thunderbolt 3, which enhanced connectivity and transfer speeds.

Ice Lake processors are available in both Core and Xeon variants. The Core lineup includes models such as the Core i3, i5, i7 and i9, offering improved performance and efficiency. The Xeon Ice Lake-SP variant, launched in 2021, features up to 40 cores, uses the BGA4189 socket and has a maximum CPU clock rate of 3.7 GHz, with performance exceeding 1 teraflop.

While the original range of Intel Core i3/i5/i7 processors from 2019 remains available, the Xeon Silver, Gold and Platinum models based on Ice Lake have been in use since 2021.

2020: Tiger Lake

The 11th-generation Intel Core mobile processors were christened Tiger Lake. They replaced the Ice Lake mobile processors, offering both dual- and quad-core models. This was the first processor since Skylake to be marketed with the Celeron, Pentium, Core and Xeon brands simultaneously.

Tiger Lake chips are the third generation of 10 nm processors and were specifically designed for lightweight gaming laptops. They offer refresh rates of 100 fps, while the Core i9-11980HK offers a maximum boost clock speed of 5 GHz.

2022: Alder Lake

Beginning in the fourth quarter of 2021 and into the start of 2022, Intel released the 12th-generation Intel mobile and desktop processor, christened Alder Lake. Alder Lake ultimately dropped the Xeon brand, although it continued to be marketed with Celeron, Pentium and Core brands simultaneously.

Using Intel's Intel 7 fabrication process, Alder Lake is described as being a 10 nm Enhanced Super Fin (ESF). The ESF was benchmarked at having 10 to 15 percent increased performance over Intel's previous 10 nm Super Fin fabrication, which was used for Tiger Lake chips. Alder Lake Core i9-12900KS offered a maximum boost clock speed of 5.5 GHz.

2023: Raptor Lake

With a limited release of certain models in the fourth quarter of 2022 and a full release in 2023, Intel released the 13th-generation mobile and desktop processor Raptor Lake. In a marked shift, Raptor Lake dropped the Celeron and Pentium processor families, instead being released with the Core and Intel Processor families.

Raptor Lake is made using the same fabrication process as Alder Lake. Raptor Lake can feature up to 24 cores and 32 threads while offering socket compatibility with Alder Lake systems. In terms of performance gains, Raptor Lake represented a half-step forward from Alder Lake, with its Core i9-13980HX processor offering a maximum boost clock speed of 5.6 GHz.

2023: Raptor Lake Refresh

Intel released a 14th-generation iterative refresh of the Raptor Lake processor for desktops at the end of 2023 and for mobile at the beginning of 2024. Known as Raptor Lake-S and Raptor Lake-XH Refresh, respectively, Intel released this generation of processors to benefit from process improvements.

Raptor Lake Refresh demonstrates substantial improvements in processing speed even over the original Raptor Lake release. The Raptor Lake Core i9-14900KS offered a maximum boost clock speed of 6.2 GHz.

2023: Meteor Lake

At the end of 2023, Intel announced the release of its first-generation Intel Core Ultra Mobile processor, Meteor Lake. Unlike past processor releases, Meteor Lake is intended just for mobile devices and will not be released for the do-it-yourself desktop PC market. It also features Intel's first use of chipset architecture.

Meteor Lake promises to be ultra-efficient while featuring various built-in functionalities. Among these are dedicated AI processors, referred to as neural processing units. Changes in Meteor Lake's architecture also mean it will have a hybrid architecture, with separate dies for the GPU, IO and system-on-chip (SOC).

Intel processor timeline

1971-81: The 4004

1978-82: iAPX 86 – 8086, 8088 and 80186 (16-bit)

1981: iAPX 432

1982: 80286

1985-94: 386 and 376

1989: 486 and i860

1993: Pentium (P5, i586)

1994-99: Bumps in the road

1995: Pentium Pro (P6, i686)

1997: Pentium II and Pentium II Xeon

1998: Celeron

1999: Pentium III and Pentium III Xeon

2000: Pentium 4

2001: Xeon, Itanium

2002: Hyper-Threading

2003: Pentium M

2005: Pentium D

2005-09: Terascale Computing Research Program

2006: Core 2 Duo

2007: Intel vPro

2008: Core i-Series, Atom

2010: HD Graphics, Many Integrated Core Architecture and Xeon Phi

2012: Intel SoCs

2013: Core-i Series – Haswell

2015: Broadwell, Skylake

2016: Kaby Lake

2017: Ice Lake

2020: Tiger Lake

2022: Alder Lake

2023: Raptor Lake

2023: Raptor Lake Refresh

2023: Meteor Lake

Intel Going Forward

Intel has had an illustrious history, helping to fuel the growth of the computing industry for both personal devices and large data centers. Recently, however, Intel has faced pressure from the rise of AMD, which has made notable gains in performance and market presence. However, despite the ongoing competition, Intel has managed to maintain a wide market share margin over AMD for personal computers. Whether that continues depends on Intel's innovation and releases in the coming years.

Jeremy Bender contributed to this article.

Did you find this content helpful?

Yes

No

Share Article:



Written by: **Neil Cumins**,
Senior Analyst

Neil Cumins is an award-winning small business owner who has run a limited company for nearly two decades. Through his personal and professional experiences, he is well-versed in a range of B2B and B2C topics, from invoicing to advertising to the use of artificial intelligence. Prior to starting his own business, Cumins worked as a marketing executive. At Business News Daily, Cumins covers technology trends, as well as specific services like GPS fleet management and contact center software. With deep insights into the ever-changing technology landscape, Cumins is particularly skilled at evaluating business software and guiding fellow entrepreneurs to the tools and strategies that will equip them for entrepreneurial success. Over the years, he has worked with some of the world's biggest hardware and software manufacturers, as well as countless SaaS brands. Today, he also spends his time consulting on compensation and other business matters.

[EMAIL](#) | [LINKEDIN](#)

Related Reviews

The Best Phone Systems for Small Businesses



The Best Business Accounting Software Services of 2025



Best Payroll Services for 2025



More Related Reviews →

Related Articles

Laptop Buying Guide for Small Business



Why You Shouldn't Cheap Out on Business Laptops



Best Networking Certifications for 2025



How to Choose the Right Multifunction Printer and Copy Machine



Best IT Certifications for 2025



More Related Articles →



info@businessnewsdaily.com
200 5th Avenue, 2nd Floor
Waltham, MA 02451



Resources >

Our Company >

Legal >

Our Brands >

© 2025 business.com All Rights Reserved.

