Futurum

intel.

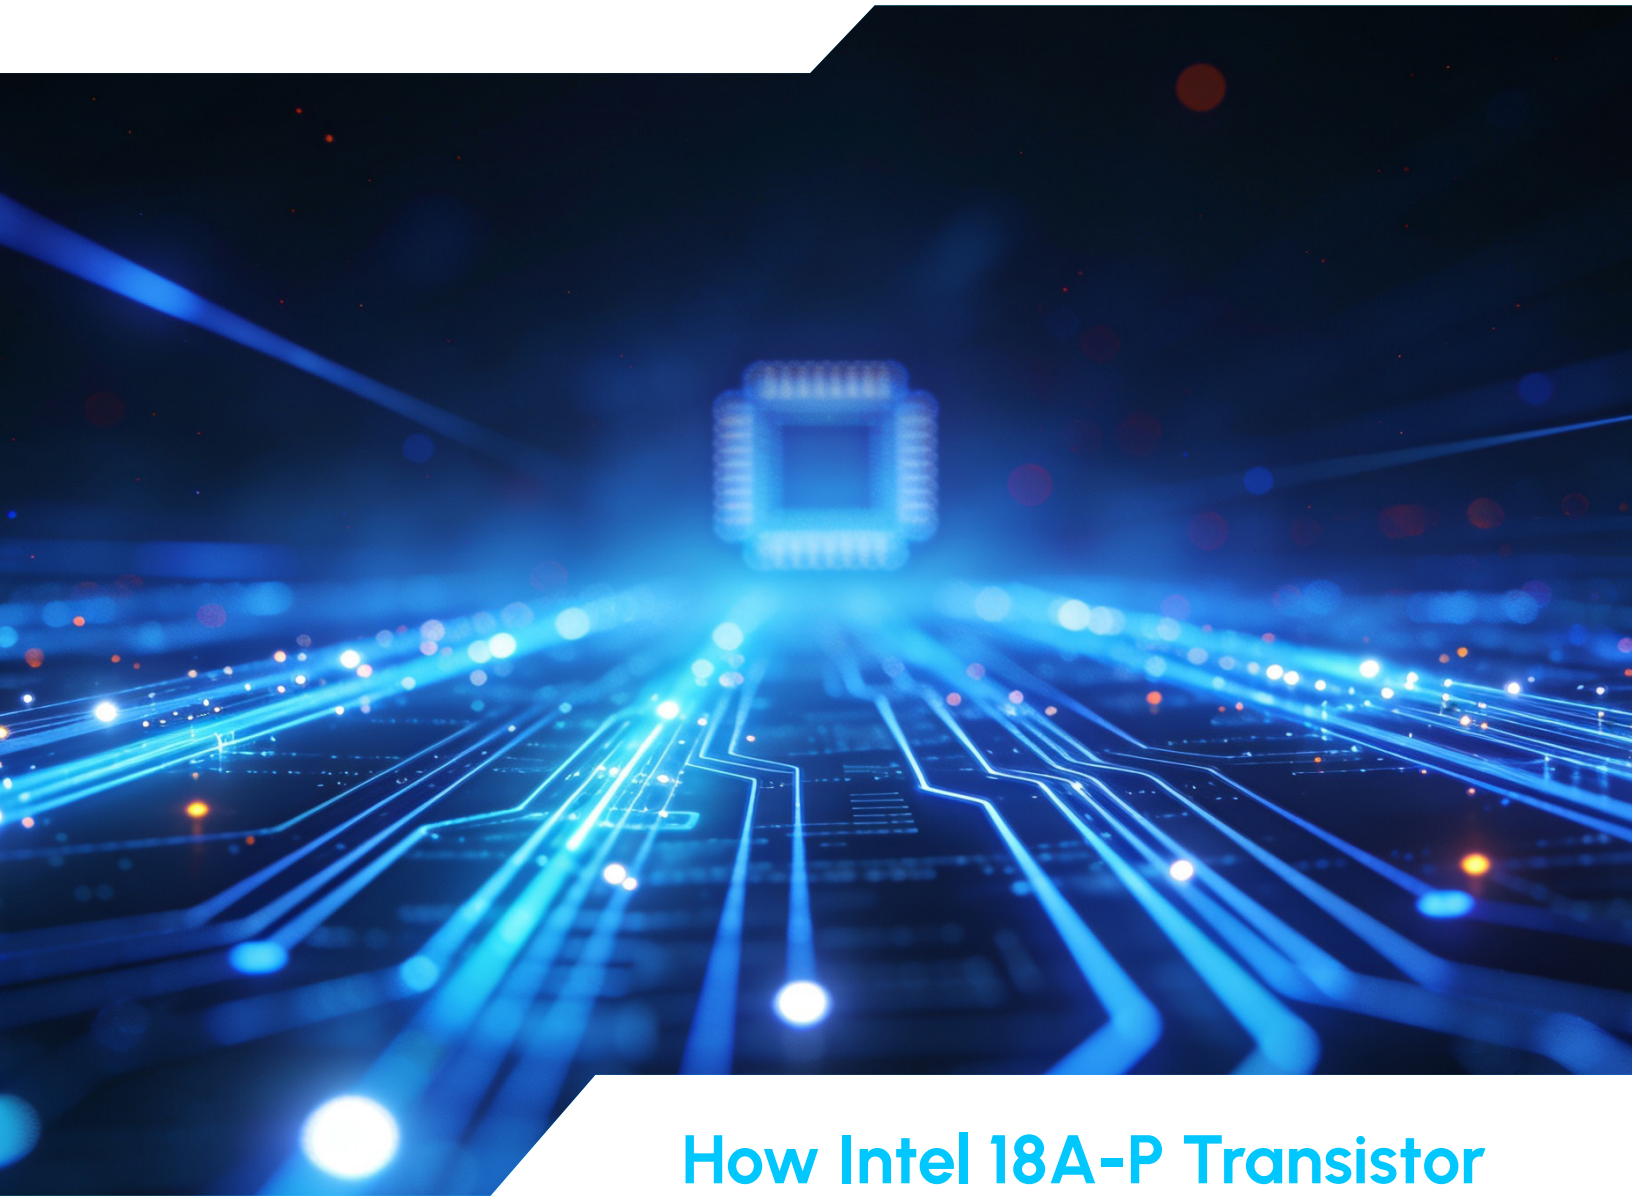

How Intel 18A-P Transistor Breakthroughs Shape CPU Design for Diamond Rapids

The first Xeon on the node builds on process gains for per-core performance

AUGUST 2026

Intel Foundry Paired Gate-All-Around With Backside Power to Buy Frequency

Intel Foundry's 18A-P process pairs gate-all-around transistors with backside power delivery (GAA-BSPD), a combination measured at about 30% higher operating frequency at low voltage on production x86 silicon. The pairing is deliberately intended to deliver gains at low voltage for efficiency and high voltage for performance. RibbonFET, Intel Foundry's gate-all-around transistor, gains most at low voltage. PowerVia, Intel's backside power delivery, routes power through the back of the wafer instead of crowding it into the wiring above the transistors, and its strengths show up at the high end. Intel's historical tick-tock cadence of density and performance gains with the prior generation of FinFET transistors sets a precedent for future Intel Foundry scaling.

Introduced in 2011, Intel's FinFET transistors drove performance gains across four process generations, from 22nm through Intel 3, each one coupling device gains to design gains (Figure 1). GAA with backside power is now at generation one, with that same runway ahead. The stakes are commercial as much as technical. Data center platforms optimize efficiency inside fixed power budgets, yet they still chase single-threaded boost frequency for the work that cannot spread across cores. The high voltage end of the curve sets that boost ceiling, and the boost ceiling is quantified in per-core performance.

Figure 1: Intel Foundry FinFET-Based Node Product History

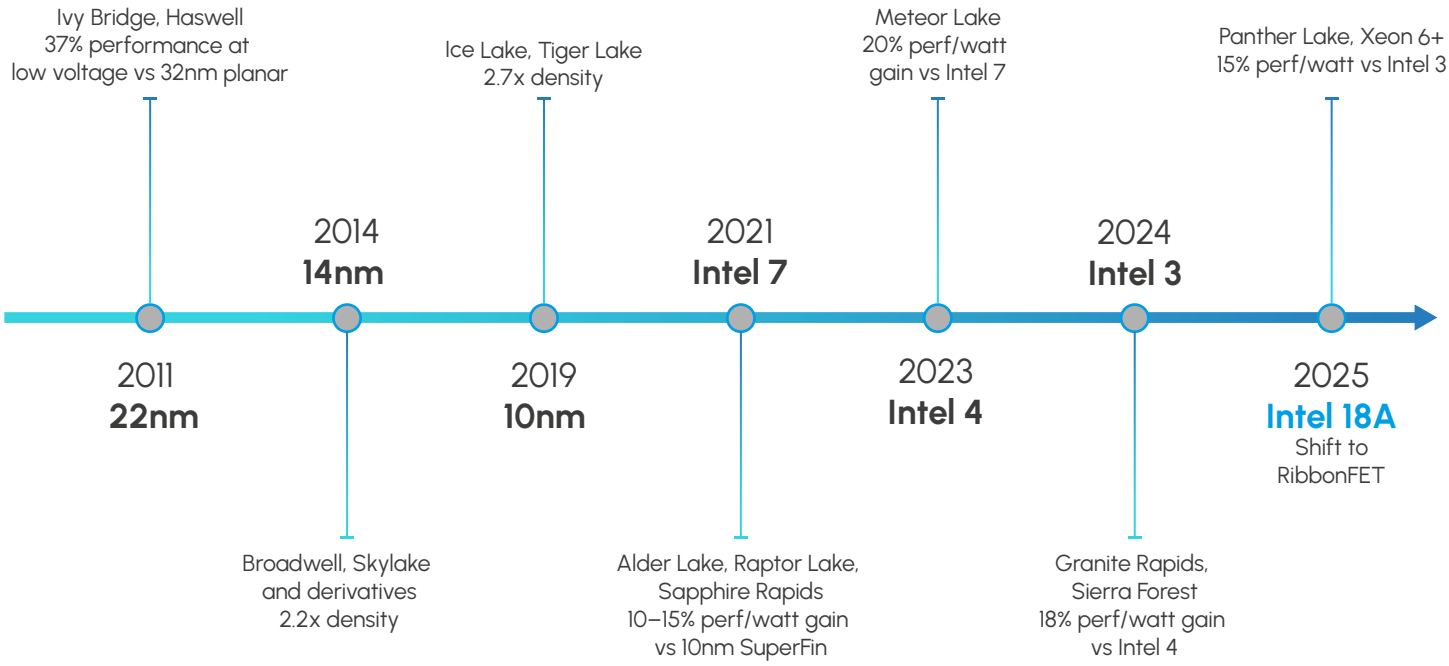

Source: Futurum Research; node years, gains, and products from Intel public disclosures[1]

1 Intel Corporation public disclosures at node and product launch, 2014–2026: node-over-node performance-per-watt claims, first-product years, and flagship products.

At Low Voltage, Intel 18A-P's Transistor Gain Translates Directly to a Product's Core

The 30% figure comes from direct measurement. Intel disclosed it in a peer-reviewed paper at the Symposium on VLSI Technology and Circuits in June 2026, the annual conference where chipmakers report measured transistor and circuit results.1 The paper presents silicon-validated design findings from CPU cores built on gate-all-around with backside power, characterized against their FinFET equivalents across the operating-voltage range. Around 0.5V, the cores reached about 30% higher operating frequency.

The comparison is taken at matched voltage on the same class of core, so it isolates the process contribution rather than an architecture change. The mechanism starts in the channel. RibbonFET's gate wraps the channel on all four sides, which gives Intel the control to lower threshold voltage. Moreover, moving power delivery to the back of the wafer cuts IR drop, the voltage lost across the delivery network before it reaches a transistor.[1]

"The whole backside power delivery concept itself really is an improvement in delivering power for high voltage. It helps at low [voltage], but at high voltage, we have large droop events where backside power is very very effective at reducing our power delivery losses. ... The pairing of backside power with gate all around is really not an accident. It's taking two things that are complementary and putting them together. So we have a story of both high and low voltage."

Eric Fetzer, Fellow, Intel[2]

Every timing path in a CPU carries two delays: the time transistors take to switch, and the time signals spend traveling the wires between them. At low voltage, the transistor runs on a thin margin between supply and threshold voltage so switching is slow and the transistor makes up almost all of the path delay. As a result, designs are device-limited at low voltage as a rule of CMOS physics. Operating in these regimes, Intel's Eric Fetzer notes, "we are typically 100% device limited in our designs. Any performance improvement in the device translates directly to frequency."[2]

What distinguishes the GAA-BSPD result is how completely that gain converts: with the device that dominant, a faster transistor becomes a faster core nearly one for one, and the measured translation exceeds what FinFET designs typically achieved. The data comes from the Core Ultra 3 Series (codenamed Panther Lake), Intel's x86 core on the Intel 18A process, captured across thousands of wafers and about 60,000 dies of high-volume manufacturing silicon.

2 Fetzer, E., & Shamanna, M. (2026, July 20). Futurum 18A-P discussion [Analyst briefing, 48 minutes; B. Fienberg, Intel, moderating; B. Burke, Futurum]. Intel Corporation.

Intel's VLSI Research Shows the Advantage Scaling to High Voltage

In their VLSI paper, "CPU Cores in GAA with Backside Power: Silicon-Validated Design Insights," the finding for high frequency designs concerns the translation of device gain into core frequency. The translation metric distinguishes how much better the transistor got and how much of that improvement shows up in clock speed. The paper finds that cores on GAA-BSPD convert front-end gains into frequency 5% better than FinFET at 1.1V and 7.5% better at 0.95V (Figure 2).[1,3] The total frequency gain at high voltage lands in double digits, before tuning from incremental design levers.

Figure 2: Share of Device Gain Reaching Core Frequency: GAA-BSPD Advantage vs FinFET

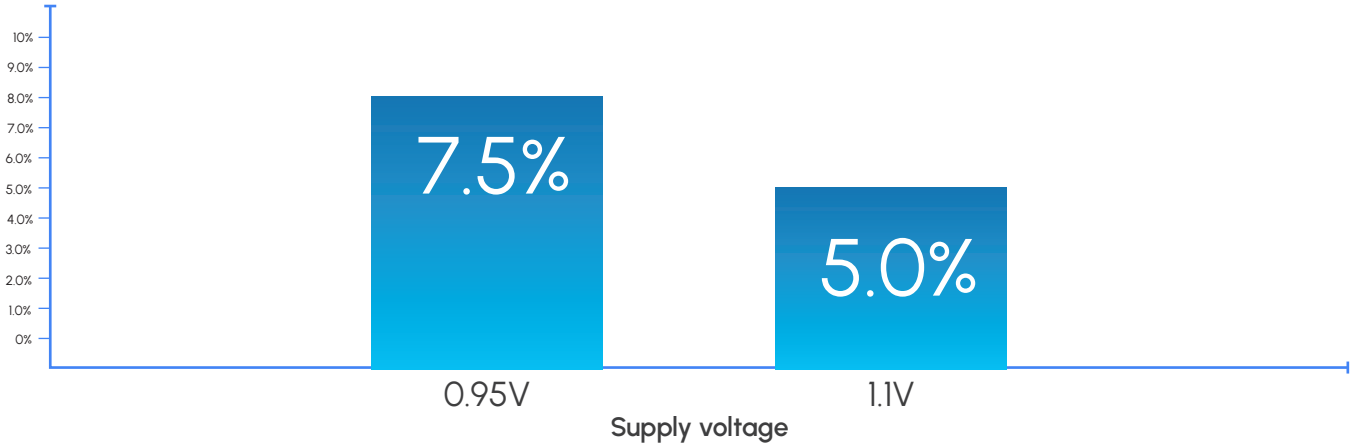

Source: Futurum Research; data from Shamanna, M., et al. (2026), Symposium on VLSI Technology and Circuits 2026, paper T2.3, Fig. 3[1]

Transistor physics explains why the low-voltage number is larger. Overdrive is the margin between supply and threshold voltage, and a transistor's drive current rises steeply with it, following a square law. Lowering threshold voltage widens that margin by the same absolute amount everywhere on the curve, but at 0.5V the addition is a large share of a small margin and at 1.1V a small share of a large one. So the same device change buys proportionally more frequency at the bottom of the curve.

The second constraint sits in the wires. At high voltage, the transistors are fast and the wires have not changed, so interconnect delay becomes an appreciable fraction of total path delay and a faster transistor no longer speeds up the core. That interconnect share is the headroom Intel's design levers now target.

3 Intel Corporation public disclosures at node and product launch, 2014-2026: node-over-node performance-per-watt claims, first-product years, and flagship products.

Four Design Levers Bring the Frequency Gain to High-Performance Designs

The high-voltage levers share a design philosophy: leave the transistor alone and clear the path around it. The most cost-effective of the four shows the performance benefits of low-cost design levers. Intel added two coarse pitch metal layers near the top of the stack, where the benefits outweigh the inexpensive costs, according to Intel's Senior Principal Engineer Manjunath Shamanna.[2] The four levers in the product Futurum sees potential to drive performance gains are demonstrated in Table 1.

Table 1: Four Levers to High-Voltage Frequency

Lever	Mechanism	Measured gain	Status
Backside power delivery (PowerVia)	Moves power to the wafer's back, suppressing droop events and freeing the front-side stack for signal routing	~10X dynamic droop reduction; 5% to 6% frequency, or more than 15% dynamic power	In production on 18A and 18A-P
Two extra coarse metal layers	Low-resistance wiring near the top of the metal stack, cutting interconnect RC and routing congestion	3% frequency at 1.1V and 2% at 0.65V, with 2% power reduction	Committed offering in 18A-P
Heat extraction through the bonding oxide	Thinner, more conductive bonding material pulls heat out at high power density, converting thermal headroom into sustained frequency	About 40% improvement across 18A generations at very high power	Delivered across 18A generations
Power boost	Front-side plus direct backside contacts, enabled by PowerVia, moving current in and out of the ribbons	Not broken out separately; part of 18A-P's device enhancement	New in 18A-P

The pattern across all four is the same. Each converts a system-level fix into clock speed, which is what translates a low voltage device advantage into the high-frequency operating points where per-core performance is decided.

Advanced Materials and Packaging Extend the Roadmap

The roadmap extends the same logic to deliver frequency gains at high voltage. Subtractive ruthenium interconnect with airgaps, demonstrated in Intel's VLSI research, cuts capacitance about 35% versus copper.[3] Its advantage compounds at fine pitch. Copper requires a resistive capping layer on all four sides that cannot shrink, and copper's resistance turns nonlinear as grain boundaries dominate at small dimensions. Ruthenium remains linear at those pitches, so it wins exactly where wire resistance limits frequency.[2]

Further out, Intel is researching logic stacked in the z dimension, one type of logic above another, so signals travel a short vertical hop instead of a long planar route. Shamanna ties that work directly to high-voltage scaling.[2] Both roadmap levers aim at the same interconnect limit the four production levers attack today, yet promise to continue gains generationally as FinFET achieved.

Diamond Rapids Converts Xeon 6+ Efficiency Into Intel's Leading High Per-Core Performance Product for Data Center

Diamond Rapids is where the process gains become a product. The next-generation Xeon is on track to deliver 50% more cores than Xeon 6 built on 18A-P performance. The trade-off behind that number started a generation earlier. Xeon 6+ ran efficiency cores, and Intel banked the resulting perf-per-watt; Diamond Rapids spends it on performance cores. As Fetzer says, "The Xeon 6+ was an efficient core versus a performance core. So we leverage that power efficiency into our performance cores for Diamond Rapids."[2]

More cores stretch the memory bandwidth that can be delivered. Intel refactored the Granite Rapids-era architecture, doubled memory bandwidth, and moved the I/O to the center of the package so every core sees uniform memory access, a layout aimed at AI-scale workloads (Table 2). The frequency headroom from 18A-P's design levers can translate to per-core performance at the system level.

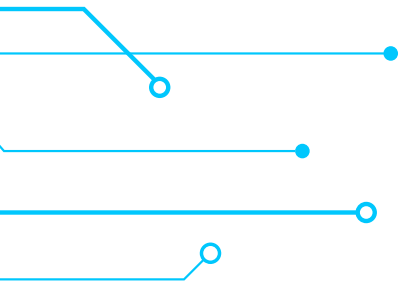

"That's what you'll see in Diamond Rapids. It starts with the underlying technology that's in the cores. We leverage the 18A-P advantages to increase both the core count and the performance of those cores, and then we update the architecture to feed it all."

Eric Fetzer, Fellow, Intel

Table 2: Diamond Rapids Product Overview

Intel 18A-P	2x memory bandwidth	PCIe Gen 6	50% more cores
Scalable SOC architecture with uniform memory latency	Increased channel count and speed for bandwidth-limited applications	Speed and scalability for I/O-heavy use cases	Same performance-core type as Xeon 6; optimized for high-demand IaaS and high performance per thread

Source: Intel

These innovations can lay the rails for a new multi-generational roadmap of per-core performance leadership. FinFET's harvest took four generations of coupled device and design work. GAA with backside power is at the first step of that curve, with a measured 30% gain at low voltage and double digits at high already on production silicon. Diamond Rapids and Nova Lake are on track to deliver the earliest of the 18A-P gains, with ruthenium interconnect and stacked logic still ahead.

Important Information About This Report

AUTHORS

Brendan Burke
Research Director, Semiconductors,
Supply Chain & Emerging Tech | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT INTEL CORPORATION

Intel (Nasdaq: INTC) designs and manufactures advanced semiconductors that connect and power the modern world. Every day, our engineers create new technologies that enhance and shape the future of computing to enable new possibilities for every customer we serve.

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

CONTACT INFORMATION: The Futurum Group LLC | futurumgroup.com | (833) 722-5337