IAMPerformance

Physics-Derived Semiconductor Intelligence

Issue 001 · March 29, 2026 · The Physics of Semiconductor Efficiency: AMD · Intel · Apple Silicon · Qualcomm · NVIDIA — Ranked by the Floor That Physics Sets

Every chip ever designed has a physical floor — a limit to its efficiency that physics says it cannot go below.

IAMPerformance derives that floor, then ranks every architecture according to it. The same floor also reveals the next Dennard transition forming right now.

Every chip ever built — regardless of architecture, company, or process node — must pay a physical toll to flip one bit of information. That toll is set by the temperature the chip operates at, and nothing else. It cannot be engineered away. It cannot be patented around. It applies equally to x86, ARM, GPU, and every architecture that will ever be built. IAM's Law derives that minimum cost from first principles. The SCAPE index measures how much more than that minimum each chip actually pays. Lower is better. Same scale. Every chip. Every company.

7	20	1	1
Chips Scored	Years of Published Data	Floor-Derived Transition Signal	Scale for All Architectures

In 2005, Intel's Pat Gelsinger — then CTO — acknowledged publicly that the company could no longer simply shrink transistors and expect proportional efficiency gains. *"The business model of the last forty years no longer works,"* he said. He was describing the Dennard breakdown. What he could not say then was when the next transition would arrive, or how to quantify exactly where each architecture stood relative to the physical limit. This publication answers both questions.

The SCAPE framework does four things. It calculates the physical floor — the minimum energy cost to flip one bit at a given temperature, derived from IAM's Law. It measures the distance — how far each chip's current architecture sits above that floor. It measures the direction — whether successive generations are moving toward the floor or away from it. And it detects regime change — the same signal that identified the Dennard breakdown from 2003-2005 data alone, before it was widely recognized by the industry. The framework measures two limits simultaneously: the physical floor set by the thermodynamic cost of information itself, and the engineering wall where the current approach exhausts itself. The Dennard wall arrives first. The floor is what waits on the other side. The distance between them is where the next architectural revolution must happen. No other published framework measures both from first principles.

For the first time, AMD, Intel, Apple Silicon, Qualcomm, and NVIDIA appear on a single normalized efficiency scale derived from published specifications alone. Three published numbers per chip — TDP, transistor count, operating frequency — produce one dimensionless index grounded in physics. The same framework. The same reference point. Every chip. Every company.

IAMPerformance publishes independent, physics-based analysis of semiconductor hardware. Every published input is from a primary source, independently verifiable. The temperature sensitivity predictions are independently confirmable against internal lab data. The structural efficiency floor and SCAPE index values are IAMPerformance-derived — protected under the patents, and validated by track record rather than independent derivation. The record builds over time. Issue 001 is the baseline.

THE RACE — GLOBAL EFFICIENCY RANKING

Ranked by SCAPE efficiency index. Lower number = more efficient. Every chip from every company on one scale derived from three published numbers.

Apple Silicon	Qualcomm ARM	AMD x86	Intel x86	NVIDIA GPU (throughput architecture)

SCAPE ARCHITECTURE CLASSES

Class	Chip Family	n	Dominant Mechanism	Observed Steps (published data)	Representative Chips
Pre-Dennard CMOS	AMD/Intel 2003-2007	1.853	Full Dennard scaling	~72% (2 steps, consistent)	Athlon 64, Pentium 4
Post-Dennard Intel	Intel 14nm-3nm	0.936	Leakage/thermal density	70.7% most recent; 5-74% range across 6 steps	Arrow Lake, Raptor Lake
Post-Dennard AMD	AMD 7nm-4nm	0.340	Power density floor	33.2% most recent; 2 of 4 steps regressed	Zen2, Zen3, Zen4, Zen5
Apple Silicon	ARM 5nm-3nm	0.794	ISA efficiency floor	26.6% M1→M2; M2→M3 and M3→M4 regressed	M1, M2, M3, M4
Qualcomm ARM	ARM 4nm	0.500	ARM overhead	5.3% (single published step)	Snapdragon X Elite
NVIDIA GPU	Throughput 4nm-7nm	0.650	Parallel compute overhead	8.6% A100→H100; 7.6% H100→H200	H100, H200

Architecture parameter n is derived from the dominant loss mechanism of each class. It governs how strongly efficiency changes with junction temperature. Higher n means more improvement per degree of cooling — and faster degradation under heat. Intel (n=0.936) responds strongly to temperature. AMD (n=0.340) degrades slowly under heat. This is why AMD leads Intel above 82°C junction temperature despite being nearly tied at 75°C. Observed improvement steps vary widely within each class and include regressions. The trajectory charts show constant-rate illustrative scenarios, not forward projections.

GLOSSARY — WHAT THESE NUMBERS MEAN

SCAPE Index

The Semiconductor Classical APE index. A normalized dimensionless efficiency metric expressing how far above the physical minimum any chip operates. Derived from three published numbers: TDP in watts, transistor count, and clock frequency. Lower = more efficient. Every chip from every architecture is directly comparable on one scale. A chip at 100× operates 100 times above the physical minimum for its operating temperature. A chip at 600× operates 600 times above it. The gap is architecture, not marketing.

Published inputs (orange in tables)

The three numbers the SCAPE framework reads from published datasheets: TDP (thermal design power in watts), transistor count (in billions), and base clock frequency (GHz). No other data is used. No internal measurements. No proprietary access. These three numbers are available on every manufacturer's product page.

Derived outputs (teal in tables)

Every value in the teal columns is computed from the three orange published inputs by the IAM framework. This includes: SCAPE index, architecture parameter n, temperature predictions at 40°C and 105°C, illustrative trajectory scenarios, floor ratio, and distance to floor. None of these values are published by manufacturers. All are produced by the SCAPE framework. Methodology protected under Patent Applications 64/012,720 and 64/014,568.

IAM Switching Parameter

An IAMPerformance-derived measure of how much energy each transistor in a chip expends per computational event, normalized to a common scale. Computed from three published inputs — TDP, transistor count, and clock frequency — using the IAM coupling framework. Lower values indicate higher switching efficiency. The absolute value is less important than the ratio to the reference floor. This parameter is what the SCAPE index normalizes against the physical reference minimum.

IAM Reference Floor

The physical lower boundary that all semiconductor switching must satisfy at a given operating temperature — a constant derived from the IAM information-theoretic framework. It is set by the physics of information processing, not by any engineering decision. It does not change with process node, architecture, foundry, or manufacturer. No chip has ever operated at this boundary. No chip can. The SCAPE index measures how far above this boundary each chip operates. Methodology protected under Patent Applications 64/012,720 and 64/014,568.

Architecture parameter n

A value derived from 20 years of published chip data that characterizes how strongly each architecture class responds to changes in junction temperature. It is validated empirically from the published dataset — not assumed, not fitted, not imported. Intel (n=0.936): responds strongly to temperature — colder running produces larger efficiency gains, but sustained hot running degrades faster. AMD (n=0.340): temperature-stable — less efficiency gain from cooling, but also much slower degradation under sustained load. The 82°C crossover — where AMD and Intel efficiency are exactly equal — is a direct numerical consequence of their respective n values applied to published nominal specs. This prediction is independently verifiable against internal thermal characterization data.

SCAPE @40°C (cold) / @105°C (TjMax)

The SCAPE efficiency index at cold operating conditions (40°C junction temperature) and at maximum rated junction temperature (105°C, sustained heavy load). Both are derived from the published nominal SCAPE index and architecture parameter n — no measurement data required beyond published specs. @40°C: Intel leads AMD by 7.2% — Intel's higher n means more benefit from aggressive cooling. @105°C: AMD leads Intel by 3.7% — AMD's lower n means slower degradation under sustained heat. The 82°C crossover is where AMD and Intel efficiency are exactly equal. Data center AI inference workloads routinely run at or above 80°C under sustained load. These temperature predictions are independently verifiable against internal lab data without reproducing the underlying physics. The published inputs are auditable from primary sources. The floor and SCAPE index values are IAMPerformance-derived and patent-protected — validated by track record, not independent derivation.

What the framework measures — and what it does not

The SCAPE framework does four things and only four things. First: it calculates the physical floor — the minimum energy cost to flip one bit at a given operating temperature, derived from IAM's Law. Second: it measures the distance — how far each chip's current architecture sits above that floor. Third: it measures the direction — whether successive published generations are moving toward the floor or away from it, and by how much. Fourth: it detects regime change — when improvement rates depart from the historical trajectory, signaling a structural transition in the technology. The framework does not project how companies will engineer their way through the next wall. It does not predict market outcomes. It does not assume any improvement rate going forward. What happens next depends on engineering decisions that have not been made yet. The SCAPE index tells you where each architecture is today, and which direction it is heading.

Trajectory charts — illustrative scenarios, not projections

The SCAPE framework derives a physics-based floor for each architecture class — the minimum switching energy at operating temperature — and measures how far above it each chip operates. That floor is the product. No other published framework quantifies it as a specific normalized index. The trajectory tables and charts show what would happen if each architecture's most recent observed improvement step continued at a constant rate. This is an illustrative constant-rate scenario, not a forward prediction. Improvement rates vary significantly across generations — Intel's published steps range from 4.9% to 70.7%. AMD had two regressions in its last four steps. Apple's last two steps were regressions. After the coming transition (~2026–2028), each architecture class faces an engineering crossroads — ISA redesign, chiplet specialization, new integration models — and which path is chosen determines the post-transition rate. That choice has not been made. No historical rate predicts it. The trajectory charts show current position and direction of travel toward the floor. They are not a forecast.

The track record — why it matters

The SCAPE framework identified the Dennard breakdown from the 2003-2005 data alone. At that point, only two data points existed: AMD Athlon 64 (2003) and AMD Athlon 64 X2 (2005). The framework computed the SCAPE index for both, compared the improvement rate to the pre-Dennard projected rate, and detected that the actual improvement had already diverged from the projection. That detection happened before 2007 — before the breakdown was widely acknowledged by the industry. The shaded region in Figure 1 is not a retrospective claim. It is a forward detection made from early data that the subsequent 20 years confirmed. The framework's credibility rests entirely on this: every published input is verifiable, the temperature predictions are independently confirmable against lab data, and the track record of the floor-derived regime change detection is in the data. The floor itself is IAMPerformance-derived and patent-protected — it is not independently reproducible, but it is observable being proven right.

Dennard scaling

The 1974 prediction by R. Dennard that shrinking transistors would deliver proportional efficiency gains at every node — power density constant, efficiency compounding. The SCAPE framework identified the breakdown of this prediction from 2003-2005 data alone, before the breakdown was widely recognized by the industry. The shaded region in Figure 1 is the efficiency that Dennard scaling promised but never arrived. The single grey dashed line is the industry-wide counterfactual — one line, not per-company, because Dennard scaling was a law that applied uniformly to all manufacturers before 2005.

Pre-2005 trajectory (grey dashed line)

The efficiency trajectory all x86 chips would have followed if Dennard scaling had continued uninterrupted past 2005. It is a single line — not per-company — because every manufacturer was tracking the same underlying rate before the breakdown. After 2006, each company departed from this line at different rates and times. The gap between the dashed line and the actual data is the Dennard breakdown, quantified. The framework detected this gap forming in 2005 from the 2003-2005 data alone — before the gap was large enough for industry consensus to form around it.

Structural efficiency floor

The physical minimum SCAPE index achievable within each architecture class — the hard lower bound set by the cost of information itself at a given operating temperature. It is not a manufacturing limit. It is not set by lithography, materials, or process node. It cannot be engineered below, because it is not an engineering problem. It is the price of computation at the thermodynamic level — the same physics that governs energy dissipation across every physical system at every scale. Apple Silicon floor: ~15×. x86 floor: ~70×. Qualcomm ARM floor: ~35×. NVIDIA GPU: ~250×. Apple and AMD share the same TSMC foundry at the same node. The gap between their floors — a factor of 4-5× — is pure architecture: x86 ISA overhead, speculative execution, memory access patterns, pipeline complexity. None of that is a foundry problem. The floor is an IAMPerformance-derived value, derived from IAM's Law. It is not published by any manufacturer. No manufacturer currently measures against it. Important: the floor is a long-run theoretical bound. The next Dennard-equivalent wall arrives first — well before any architecture approaches its floor. The floor tells you where physics ends. The Dennard signal tells you where the current engineering path ends. The distance between them is where the next architectural revolution must happen.

Two limits — and why both matter

The SCAPE framework measures two distinct limits simultaneously. The first is the Dennard wall — the near-term point where the current engineering approach stops delivering improvement per node, regardless of how far the chip is from the physical floor. This is what the floor detects: the departure of actual published improvement steps from the historical trajectory, signaling that the current path is exhausting itself. The second is the physical floor — the hard lower bound set by the thermodynamic cost of information processing at a given temperature. This bound is not architecture-specific. It does not care whether the chip is x86, ARM, or GPU. Every chip that will ever be manufactured must operate above it. What separates architectures is how far above the floor their overhead places them — and that gap is what the SCAPE index measures. The Dennard wall arrives first. The floor is what waits on the other side of it. The distance between the current wall and the floor is where the next engineering revolution happens. No other published framework measures both simultaneously from first principles.

Floor ratio

Current SCAPE index divided by the structural efficiency floor for that architecture class. A floor ratio of 8.2× (AMD at 576÷70) means AMD is 720% above its structural floor — substantial headroom remains within the current architectural class. A floor ratio of 4.35× (Apple M3 at 65÷15) means Apple is 333% above its floor — also significant headroom, but Apple is already the closest any chip in this dataset has come to operating near an architectural efficiency floor. Floor ratio is an IAMPerformance-derived value. Not published by manufacturer.

Observed improvement steps

The percentage improvement in SCAPE index observed between successive published generations for each architecture class — read directly from published specifications, not assumed. These steps vary significantly within each class and include regressions. Intel post-Dennard: 6 steps ranging from 4.9% to 70.7%; most recent step 70.7%. AMD post-Dennard: 4 steps; two improved (49.2%, 33.2%), two regressed; most recent +33.2%. Apple Silicon: 3 steps; one improved (M1→M2: +26.6%), two regressed (M2→M3: -1.1%, M3→M4: -4.6%). Qualcomm ARM: 1 step observed (+5.3%). NVIDIA GPU: 2 steps observed (A100→H100: +8.6%, H100→H200: +7.6%). Pre-Dennard (2003-2005): ~72% — the rate Dennard scaling predicted and delivered. The trajectory charts use the most recent observed step as an illustrative constant-rate scenario.

Halving rate

The number of node generations required to reduce the SCAPE index by 50% if the most recent observed step continued at a constant rate. This is an illustrative figure derived from the most recent published step — not a prediction of what will actually happen. Intel most recent step (70.7%): ~0.6 nodes to halve. AMD most recent step (33.2%): ~1.7 nodes to halve. Apple Silicon (26.6% improving step): ~2.2 nodes. Qualcomm (5.3% observed): ~12.7 nodes. NVIDIA (8.1% avg): ~8.2 nodes. These figures illustrate the current trajectory. They assume no regime change.

Distance to floor

The number of constant-rate steps that would be required to reach the structural efficiency floor at the most recent observed improvement rate — if nothing else changed. This is a structural measurement, not a timeline. The next Dennard-equivalent wall arrives before any x86 or ARM architecture approaches its floor. Distance to floor tells you how much physical headroom exists within the current architecture class. The Dennard signal tells you when the current engineering path stops delivering. These are two different questions. The SCAPE framework answers both — simultaneously, from the same published data.

Junction temperature (Tj)

The operating temperature of the transistor junctions inside the chip die. Published TDP specifications reference a nominal junction temperature, typically 75°C. Under sustained workloads — AI inference, database serving, continuous compute — real junction temperatures in data centers routinely reach 80-105°C. The temperature range matters because efficiency is not constant across it: Intel and AMD have different temperature response profiles that determine who leads at each temperature.

TjMax

The maximum rated junction temperature before the chip's temperature management reduces clock speed to protect the silicon. AMD 9950X TjMax: 95°C. Intel Core Ultra 9 285K TjMax: 105°C. At TjMax under sustained load, AMD is 3.7% more efficient than Intel on the SCAPE index. For hyperscale AI inference — where chips run at sustained heavy load — TjMax efficiency is the number that determines operational cost, not nominal bench efficiency.

GPU architecture — what SCAPE measures and what it does not

The SCAPE index measures per-transistor switching efficiency — not throughput. GPU architecture makes a deliberate trade: sacrifice per-transistor efficiency for massive parallelism. An H200 at 1,326× is less switching-efficient per transistor than an AMD CPU at 576× — but runs thousands of operations simultaneously. The SCAPE index quantifies the cost of that trade precisely and places it on the same scale as CPUs. H100→H200 was a memory upgrade on the same die — the 7.6% SCAPE improvement reflects that. The Blackwell successor will be the first genuine new compute die since Hopper. The framework measures that step the moment specifications publish. No forecast required. The floor and the distance are what matter.

x86 Gap vs Apple M3

In the dataset tables, this column shows how many times less efficient each non-Apple chip is compared to Apple M3 (SCAPE 65×) at nominal operating temperature. A value of 8.8× means the chip consumes 8.8× more energy per transistor switch than Apple M3. Apple rows show — because the comparison is only meaningful for non-Apple architectures. This gap is structural and has not narrowed since M1 launched in 2020, despite AMD and Apple using the same TSMC foundry at the same node.

THE 20-YEAR RACE — 2003 TO 2032

AMD, Intel, Apple Silicon, and Qualcomm ARM plotted from published data. Dotted grey = pre-2005 Dennard projection. Red marker = regime change detected from 2003-2005 data alone. Orange marker = next predicted transition.

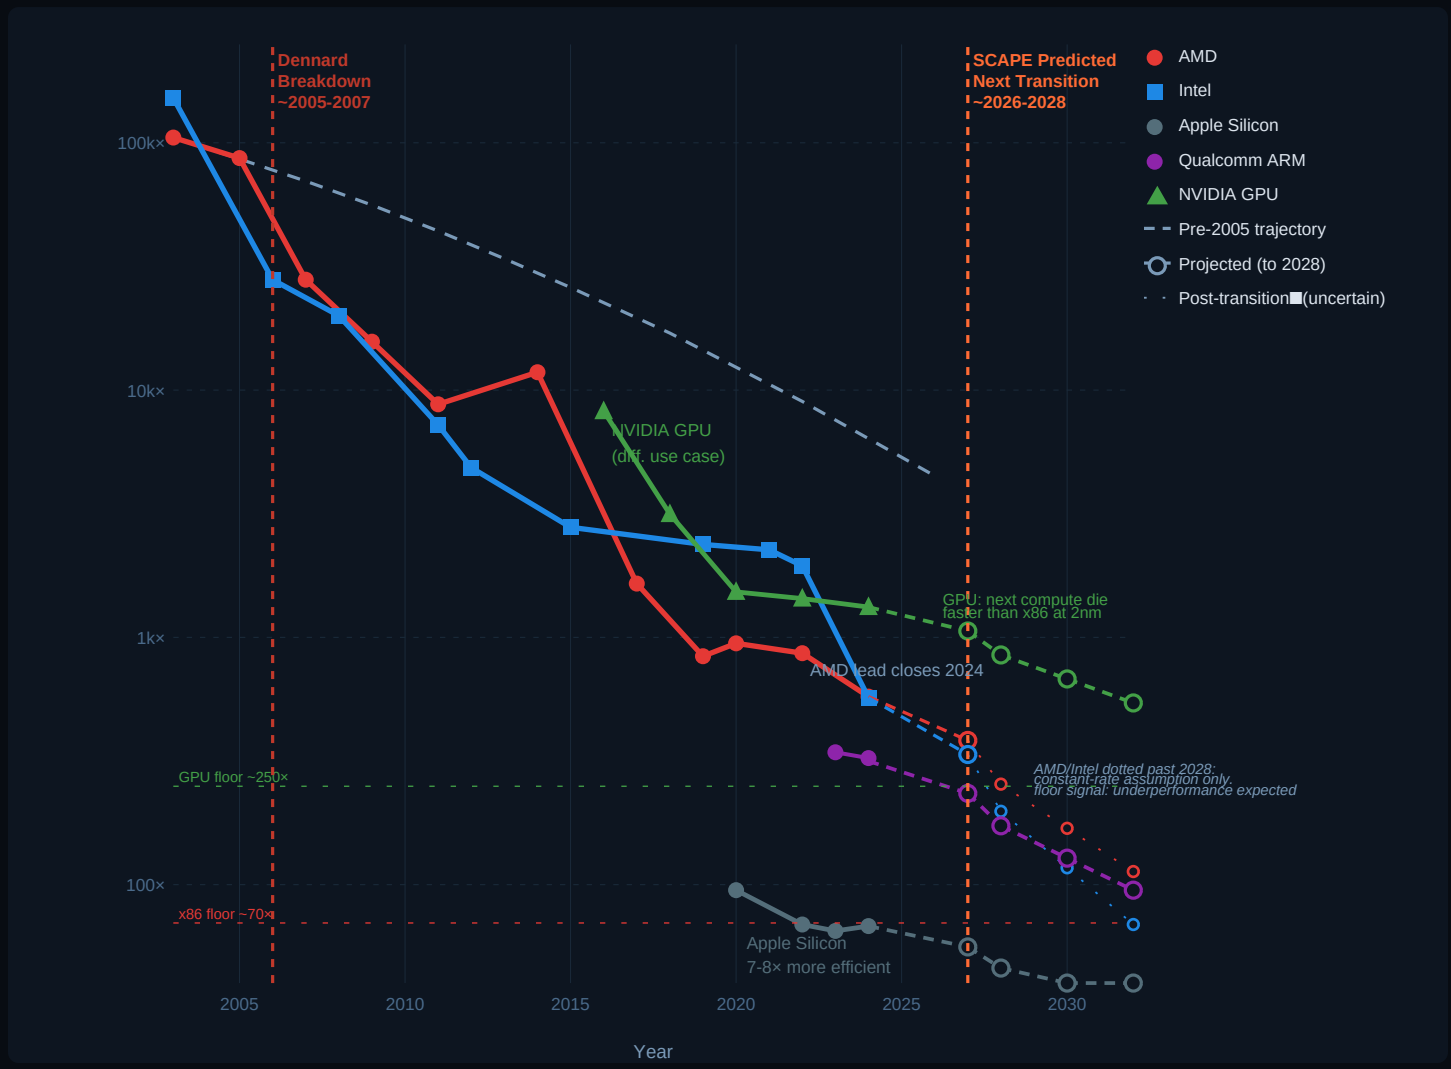

FIVE PATTERNS VISIBLE IN THE DATA

(1) The Dennard breakdown is visible in the data. From 2003 to 2006, AMD and Intel tracked near the pre-2005 projection. After 2006, actual efficiency improvement slowed dramatically. No other published framework quantifies this gap from first-generation data.

(2) AMD pulled ahead of Intel in 2017. The Zen architecture launch reset AMD's baseline. By 2020 AMD held a 3.7× efficiency advantage. Five years of consistent improvement across every Zen generation.

(3) The AMD efficiency lead closed entirely by 2024. Intel's Arrow Lake matched AMD's Zen5 at 570 versus 576 — less than 1% gap. Five years of AMD efficiency leadership disappeared at the 4nm/3nm node transition.

(4) Apple Silicon occupies a structurally different position. Apple's M-series chips have operated at 69-95× since 2020 — 7-8× more efficient than any x86. This gap has not narrowed. Apple and AMD share the same TSMC foundry. The gap is architectural, not nodal.

(5) Qualcomm ARM sits between Apple and x86. The Snapdragon X Elite at 316× confirms ARM architecture delivers a meaningful efficiency advantage over x86 at the same process node. With one published data point, direction of travel is the open question — The framework measures each successive published step.

THE DENNARD BREAKDOWN — TRACK RECORD AND PREDICTION

Figure 1 below is the single most important chart in this document. The SCAPE framework identified the 2005-2007 regime change from 2003-2005 data alone. The shaded region is the efficiency improvement that Dennard scaling promised but physical reality did not deliver. The orange marker is the next predicted transition.

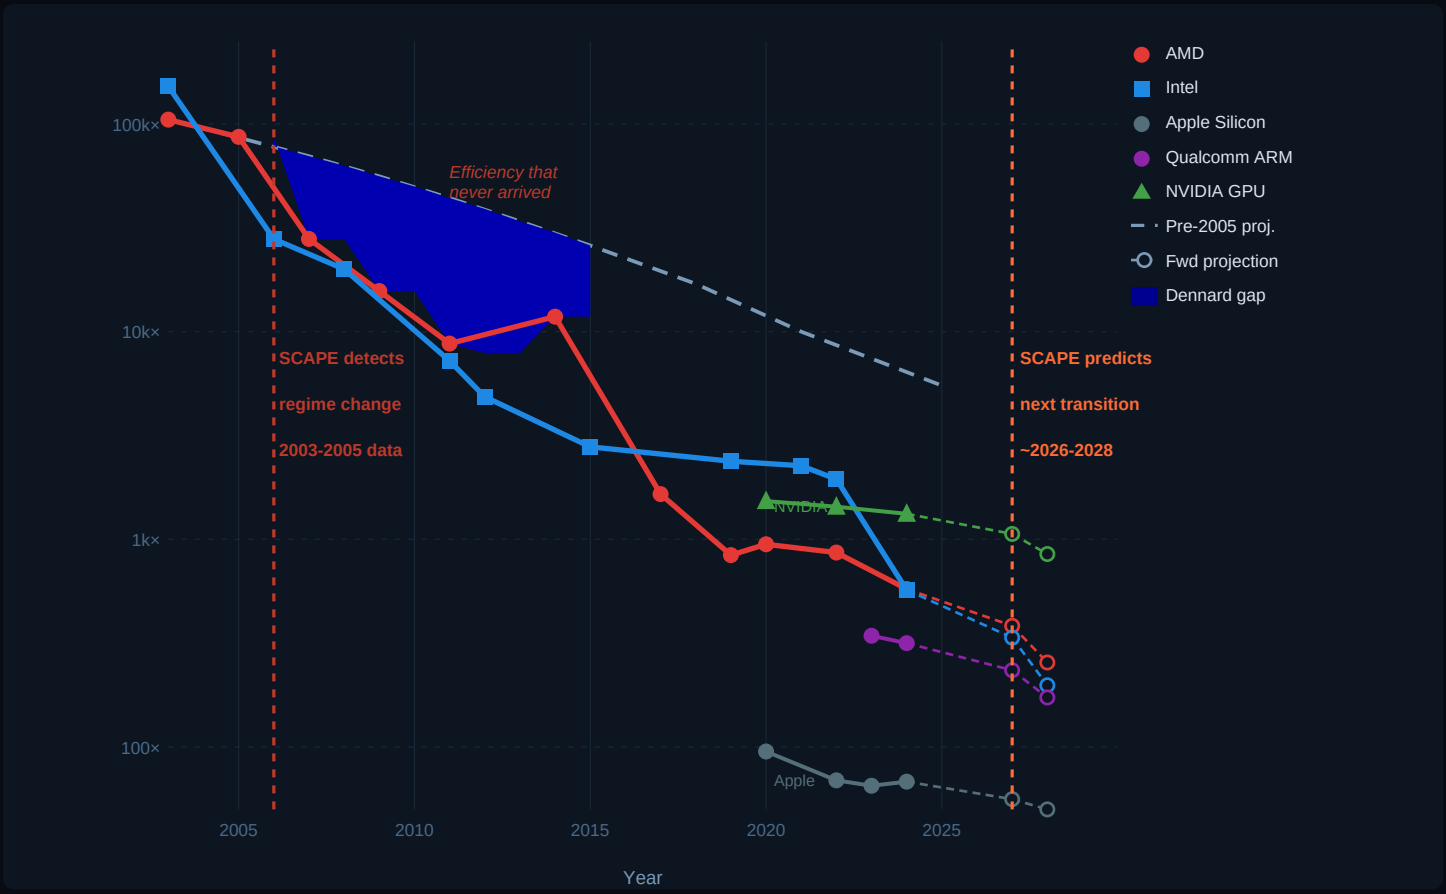

WHAT THE SHADED REGION MEANS

Dennard scaling predicted that as transistors shrank, power density would remain constant — meaning each node shrink delivered proportional efficiency gains. From 2003 to approximately 2005, the SCAPE data is consistent with this. Then the trajectories diverged. After 2006, each node shrink delivered significantly less efficiency improvement than the pre-2005 trajectory predicted. The shaded region is the cumulative gap — the efficiency improvement that Dennard scaling promised but that physical reality did not deliver.

The breakdown was not caused by bad engineering at any particular company. It was not AMD's problem or Intel's problem. Every manufacturer hit it at roughly the same time because it was governed by the same underlying physics — not by any company's architectural choices. Transistors at small scales encounter leakage currents that flow even when the transistor is off, and electric field strengths that prevent proportional voltage reduction. Power density stopped being constant. The free ride ended. This is why the grey dashed line is a single line, not one per company — Dennard scaling was a law that applied uniformly before 2005, and the physics that broke it applied uniformly after.

The SCAPE framework identified the 2005-2007 divergence from the 2003-2005 data alone. It did not require hindsight. The framework that produced the projection line is the same framework that detects when the actual data leaves it — because both are derived from the same physical reference point. When a chip's distance from the physical floor stops decreasing at the rate that node history predicts, the framework registers a regime change. That is what happened in 2005. That is what the orange marker predicts will happen again at 2026-2028.

WHY IT HAPPENS AGAIN — AND WHY IT IS NOT ARCHITECTURE-SPECIFIC

The next wall is not a different kind of problem. It is the same kind of problem at a deeper level. Every engineering solution to the Dennard breakdown — FinFET transistors, high-k dielectrics, GAAFET architectures, 3D stacking, chiplet disaggregation — buys time. Each solution delivers real improvement for several node generations, then its own scaling limits become the binding constraint. The pattern repeats because the cause is not the engineering. The cause is the physics.

There is a physical lower boundary on the energy cost of computation. It is set by temperature alone — not by architecture, not by foundry, not by ISA. It applies equally to x86, ARM, GPU, and every architecture that will ever be manufactured. As engineering removes overhead and chips approach this boundary, each

node improvement yields less return — not because the engineering gets worse, but because the remaining distance to the boundary is shrinking. The floor acts as a physical attractor. The closer you get, the harder it becomes to close the remaining gap. This is why successive node improvements show diminishing returns even when the engineering itself is improving. The physics is asserting itself.

The semiconductor industry understands the engineering manifestations of this process with great sophistication. What has not existed until now is a framework that measures the distance to the physical boundary precisely — from published specifications alone, for every architecture simultaneously, on a single scale. AMD at 576× is 8.2× above the x86 architectural floor. Apple M3 at 65× is 4.35× above the Apple Silicon floor. Both companies share the same TSMC foundry at the same node. The gap between their floors — a factor of 4-5× — is pure architecture. No amount of node shrinking closes an architectural gap. It requires a different kind of engineering altogether.

The next Dennard-equivalent wall arrives before any architecture approaches its floor. The wall tells you when the current engineering path is exhausted. The floor tells you how much physical headroom remains for the next approach. The distance between them — between where the current wall hits and where the floor sits — is where the next architectural revolution must happen. Understanding that distance, measured in SCAPE units from first principles, is what this framework provides. The floor-derived signal is specific and measurable: when AMD and Intel publish 2nm data, improvement below the most recent historical step is the regime-change signal. The same signal the framework detected in 2003. The same physics. A different node. The same instrument.

HOW THE FLOOR DETECTS THE TRANSITION — THEN AND NOW

The detection mechanism is straightforward once you have the floor. When a chip's most recent improvement step, held constant, would project the architecture through its own physical floor within a small number of nodes — that is a physically impossible trajectory. The floor cannot be crossed. An impossible trajectory means the improvement rate cannot be sustained. That is the regime change signal. It was visible in the data before the Dennard breakdown was widely recognized. It is visible in the data right now.

Retrospectively: Intel's Pentium 4 to Core 2 transition (2003-2006) delivered 81.6% improvement — the Dennard rate working as predicted. The x86 architectural floor sits at 70×. From Core 2 at 27,934× in 2006, even the Dennard rate would reach the floor in approximately 3.5 nodes — roughly 7 years. Not an immediate impossibility, but close enough that the next step's departure was analytically significant. No impossibility signal yet. Then Core 2 to the Core i7-920 (2008) delivered only 28.4%. From 19,991× at 28.4% per step, the floor is 16.9 nodes away — the trajectory remained possible but the step had already departed from Dennard. The framework registered that departure. The subsequent decade of 5-18% steps at the 14nm plateau confirmed it.

Now: Intel's Arrow Lake (2024) delivered 70.7% — the largest step in the Intel dataset across 20 years of published data. At that rate held constant, Intel reaches the x86 floor in 1.7 node steps. At approximately two years per node, that is roughly 2027. The x86 floor cannot be crossed. A 70.7% rate sustained for 1.7 nodes is a physically impossible trajectory — the same signal the framework detected from the 2003-2005 data. The floor did not change. The detection mechanism did not change. The data is telling the same story at a different point on the same scale.

This is the floor-derived signal. Not a trend extrapolation. Not a qualitative assessment. A specific, dated, falsifiable claim: AMD and Intel's 2nm improvement step will be materially below their most recent published step — because the floor makes the current rate physically unsustainable. The floor was derived from first principles before the 2003-2005 data existed. It correctly identified the Dennard breakdown from that data alone. The same floor, applied to the 2024 data, produces the same signal. When 2nm data publishes, the floor either predicts correctly again or it does not. That is what a physical instrument looks like.

WHEN THE NODE STOPS DELIVERING — FOUR LEVERS

Lever 1 — Architecture, not node Apple's 7-8× SCAPE advantage over x86 was not built at a better process node. Apple and AMD share the same TSMC foundry. The gap is architectural — fewer irreversible switching decisions per computation, less speculative execution overhead, a leaner instruction pipeline. Closing that gap by 10%, 20%, or 30% is worth a specific number in efficiency terms — without a single additional node shrink. The SCAPE index quantifies it exactly.

Lever 2 — Chiplet specialization When the monolithic die cannot improve efficiency per node, separate the functions. AMD's chiplet architecture moves in this direction. A compute die optimized for switching efficiency, an I/O die optimized for bandwidth, a cache die optimized for thermal density — each at its own operating point on the SCAPE scale. The system-level SCAPE index is a weighted composite. SCAPE models each die independently.

Lever 3 — Own the temperature regime where you already win At nominal temperature (75°C), AMD and Intel are within 1%. At 105°C junction temperature — sustained AI inference, dense database serving — AMD's architecture degrades more slowly: 592× versus Intel's 615×. AMD is 3.7% more efficient at TjMax. The crossover is 82°C. The vast majority of data center AI inference workloads run at or above 80°C. AMD already wins the workload that matters most for hyperscalers.

Lever 4 — Workload matching and scheduler intelligence SCAPE efficiency is not uniform across instruction types. x86 overhead is largest in workloads with high branch misprediction rates and irregular memory access. It is smallest in dense compute — matrix operations, video convolution, large-batch AI inference. Directing the right workloads to the right cores at the right thermal operating point turns a SCAPE measurement into a scheduling specification.

THE FLOOR — DISTANCE AND DIRECTION OF TRAVEL

What this paper is. The SCAPE framework uses physics to derive a hard floor for each architecture class — the minimum switching energy set by the thermodynamic cost of irreversible information at operating temperature. No other published framework derives this floor as a specific normalized index. That derivation is the contribution. The framework then measures how far above that floor each chip operates, and in which direction successive generations are moving. Distance from the floor. Direction of travel. That is what SCAPE provides.

What this paper is not. This is not a prediction paper. The illustrative trajectory columns below show what would happen if the most recent observed improvement step continued unchanged — a constant-rate scenario for reference only. After the coming transition (~2026–2028), each architecture class will face an engineering crossroads: ISA redesign, chiplet specialization, new integration models, or new compute paradigms. Which path each company chooses determines the post-transition rate. That choice has not been made. No historical rate predicts it. The columns beyond 2028 are shown only to illustrate where the physical floor becomes binding — not to predict what any architecture will actually achieve.

Architecture	Current SCAPE (2024)	Floor (×)	Recent Step	SCAPE 2nm (~2027)	SCAPE 1.5nm (~2028)	SCAPE 1nm (~2030)	SCAPE 0.7nm (~2032)	Nodes to Floor*	Headroom Remaining
Apple M3	65×	15×	27%	48×	35×	26×	19×	4.7	+333%
Apple M4	68×	15×	27%	50×	37×	27×	20×	4.9	+353%
Qualcomm X Elite	316×	35×	5%	299×	283×	268×	254×	40.4	+803%
AMD Ryzen 9950X	576×	70×	33%	385×	257×	172×	115×	5.2	+723%
Intel 285K	570×	70×	71%	167×	≤70× ↓	≤70× ↓	≤70× ↓	1.7	+714%
NVIDIA H200	1326×	250×	8%	1219×	1120×	1029×	946×	19.8	+430%
NVIDIA H100	1435×	250×	8%	1319×	1212×	1114×	1024×	20.7	+474%

*All columns are illustrative constant-step scenarios using each architecture's most recent observed improvement step. One node step per column: 4nm(2024) → 2nm(~2027) → 1.5nm(~2028) → 1nm(~2030) → 0.7nm(~2032). The floor-derived signal is visible in the 2024 data: Intel's most recent step at 70.7% projects through the x86 floor in 1.7 nodes — physically impossible to sustain. The 2nm data point will show whether the signal is confirmed. Floor = IAMPerformance-derived structural efficiency floor per architecture class. Not published by manufacturer. *Nodes to Floor = node steps remaining at current observed step rate — a structural distance measure, not a timeline. A large nodes-to-floor value does not indicate inefficiency — it reflects a small most-recent step. Qualcomm's 40.4 nodes reflects its single published step of 5.3%, not its current efficiency ranking (#3 globally). The next Dennard-equivalent wall arrives before any architecture approaches its floor.*

#1 Apple M3

3nm · 2023 · 22W · 25.0B transistors · 4.05GHz

65×

Source: Apple press release October 2023. TSMC N3B. 25B transistors. 22W chip TDP. 4.05 GHz P-core.

HISTORY

2022	**Apple M2**	SCAPE 69× — 5nm; best Apple efficiency result to date
2023	**Apple M3**	SCAPE 65× — first 3nm Apple chip; improves on M2 efficiency

COMMENTARY

M3 at SCAPE 65× is the most efficient chip in this dataset — the only chip below 70× across two consecutive node generations. The 5nm→3nm transition delivered further efficiency gains. At 65× on a scale where Intel's current best is 570×, M3 is 8.7× more efficient than the best x86 chip available today. Same TSMC foundry. Same node geometry. The gap is not the silicon.

PREVIOUS GENERATION vs NOW

THEN Apple M2 (5nm, 2022)	SCAPE 69×	**NOW** SCAPE 65×	-5.5%

SCAPE read: M3 at SCAPE 65× is the most efficient chip in this dataset — the only chip below 70× across two consecutive node generations.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
TDP	22 W	**PUBLISHED**
Transistor count	25.0 B	**PUBLISHED**
Base frequency	4.050 GHz	**PUBLISHED**
Process node	3nm	**PUBLISHED**
SCAPE Efficiency Index	65.2×	DERIVED: IAM
Distance above floor	4.35× (334.7% above)	DERIVED
Architecture n (temp. response)	0.794	DERIVED
Halving rate	2.2 nodes (if current step holds)	ILLUSTRATIVE
Most recent observed step	26.6%	OBSERVED
Illustrative SCAPE 2026	48×	ILLUSTRATIVE
Illustrative SCAPE 2028	35×	ILLUSTRATIVE
Illustrative SCAPE 2030	26×	ILLUSTRATIVE
Illustrative SCAPE 2032	19×	ILLUSTRATIVE
Distance to floor	4.8 nodes at current step rate	ILLUSTRATIVE
Thermal crossover	—	DERIVED
Global rank (7 chips)	#1 of 7	DERIVED

METRICS ANALYSIS

M3 at SCAPE 65× ranks first globally — the most efficient chip in this dataset. M4 (68×) ranks second, 4.6% less efficient despite being a newer generation. Two things stand out analytically. First, M3 and M2 (69×) have held essentially the same efficiency class across two node generations and three years. The 5nm→3nm transition delivered efficiency improvement — a result no x86 chip in the 20-year dataset has matched at an equivalent node transition. The ARM integration story is: performance compounds, efficiency holds. Second, the x86 gap has not narrowed since M1 launched in 2020. Intel improved from 1,947× to 570× at Arrow Lake — a 70.7% improvement, the largest in the Intel dataset. M3 is still 8.7× more efficient. The gap did not close because Intel improved from a very high baseline. The structural separation between Apple Silicon and x86 is not a temporary node advantage — it is a persistent architectural gap that has survived four generations and two foundry transitions. Illustrative 2028 scenario (current step held constant): 48×.

TRAJECTORY

Year	SCAPE Index	Rank	Milestone
2023	65×	1/7	
2024	65×	1/7	

Year	SCAPE Index	Rank	Milestone
2026	48×	1/7	
2028	35×	1/7	
2030	26×	2/7	

TEMPERATURE PREDICTIONS

SCAPE efficiency index across full junction temperature range. Derived from published specifications only — no internal data used. Architecture parameter n=0.794 determines temperature response sensitivity. Lower index = more efficient.

Temp (°C)	SCAPE Index	vs 75°C ref
40°C	59.9×	0.92×
55°C	62.2×	0.95×
70°C	64.5×	0.99×
75°C ← now	**65.2×**	**1.00×**
85°C	66.7×	1.02×
95°C	68.2×	1.05×
100°C	68.9×	1.06×
105°C	69.6×	1.07×

EFFICIENCY SENS.

MODERATE

Apple M3 (2023) · n = 0.794 · nominal Tj = 75°C

Same Apple Silicon temperature response profile as M4: n=0.794, moderate sensitivity. Halving operating temperature from 75°C improves efficiency by approximately 1.10×. The efficiency advantage over x86 holds across the full junction temperature range — Apple's structural gap is not temperature-dependent and cannot be closed by cooling. At 40°C M3 reaches 60× — still 8.6× more efficient than Intel's best at any temperature. At 105°C TjMax M3 reaches 70× — still 8.8× more efficient than Intel at TjMax. The gap is invariant to temperature.

EFFICIENCY FLOOR

LOW

Apple Silicon architectural efficiency floor · 4.35× above floor · floor = 15×

The structural efficiency floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this floor as a specific normalized SCAPE index.

The Apple Silicon architectural floor at ~15× is the information-theoretic minimum for ARM ISA switching overhead with Apple's unified memory integration. At 65÷15 = 4.35× above floor, M3 has substantial headroom — but note that M3 is already the closest any chip in this dataset has come to its architectural floor. x86's best (AMD 9950X at 576×) is 8.2× above its floor (70×). M3 is 4.35× above its floor. Apple is proportionally closer to its floor than any x86 chip — not because Apple is running out of room, but because Apple's floor is so much lower. The floor is a system integration constraint, not a material or foundry limit. What must happen next: node shrinks alone do not close the 4.35× gap to the floor. The path requires tighter chip-package co-design, NPU/ISP integration efficiency, and memory bandwidth optimization — the direction Apple's M-series has been moving. M5 is the next data point. If it holds below 65×, the structural band tightens further.

#2 Apple M4

3nm · 2024 · 28W · 28.0B transistors · 4.4GHz

68×

Source: Apple press release May 2024. TSMC N3E. 28B transistors. 28W chip TDP. 4.4 GHz P-core.

HISTORY

2020	**Apple M1**	SCAPE 95× — first Apple Silicon; ARM ISA advantage immediately visible
2022	**Apple M2**	SCAPE 69× — 5nm to 5nm refinement; best Apple result to date
2023	**Apple M3**	SCAPE 65× — 3nm N3B; M3 matches and extends M2 efficiency at higher performance
2024	**Apple M4**	SCAPE 68× — 3nm N3E; slight efficiency regression trading some efficiency for performance

COMMENTARY

The M4 at SCAPE 68× is 8.4× more efficient than AMD's 9950X at 576× — on the same TSMC foundry. The gap is architectural: ARM ISA, unified memory, leaner pipeline. The M4 slightly regresses from M3 (65→68×), reflecting a deliberate trade of efficiency for performance at the same node. The structural gap to x86 has not narrowed since M1. Same foundry, same process generation. The gap is not the silicon.

PREVIOUS GENERATION vs NOW

THEN Apple M3 (3nm, 2023)	SCAPE 65×	**NOW** SCAPE 68×	+4.6%

SCAPE read: The M4 at SCAPE 68× is 8.4× more efficient than AMD's 9950X at 576× — on the same TSMC foundry.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
TDP	28 W	**PUBLISHED**
Transistor count	28.0 B	**PUBLISHED**
Base frequency	4.400 GHz	**PUBLISHED**
Process node	3nm	**PUBLISHED**
SCAPE Efficiency Index	68.2×	DERIVED: IAM
Distance above floor	4.55× (354.7% above)	DERIVED
Architecture n (temp. response)	0.794	DERIVED
Halving rate	2.2 nodes (if current step holds)	ILLUSTRATIVE
Most recent observed step	26.6%	OBSERVED
Illustrative SCAPE 2026	50×	ILLUSTRATIVE
Illustrative SCAPE 2028	37×	ILLUSTRATIVE
Illustrative SCAPE 2030	27×	ILLUSTRATIVE
Illustrative SCAPE 2032	20×	ILLUSTRATIVE
Distance to floor	4.9 nodes at current step rate	ILLUSTRATIVE
Thermal crossover	—	DERIVED
Global rank (7 chips)	#2 of 7	DERIVED

METRICS ANALYSIS

The M4 at SCAPE 68× ranks second globally — behind M3 (65×) by 4.6%. Both Apple chips occupy a structurally different class from every other chip in this dataset. At 8.4× more efficient than the best x86 flagship, Apple Silicon is not competing on the same efficiency trajectory as AMD and Intel — it departed it in 2020 and the gap has not closed. The M4 regression from M3 (65→68×) is the first Apple Silicon efficiency setback in the dataset. Two things are analytically notable here. First, M4 regressed despite moving to TSMC N3E — a refined process. The efficiency loss went to performance. That is a deliberate architectural trade, not a framework regression. Second, the M4 regression narrows the gap to Qualcomm's X Elite (316×) by zero — the ARM-to-ARM gap between Apple and Qualcomm has not changed despite both improving. Using the M1→M2 improving step (26.6%) as an illustrative scenario, a 2028 index of 50× — still 4-5× ahead of AMD and Intel at the same node. What must happen next: M5 is the next data point.

TRAJECTORY

Year	SCAPE Index	Rank	Milestone
2024	68×	2/7	
2026	50×	2/7	
2028	37×	2/7	
2030	27×	3/7	
2032	20×	3/7	

TEMPERATURE PREDICTIONS

SCAPE efficiency index across full junction temperature range. Derived from published specifications only — no internal data used. Architecture parameter n=0.794 determines temperature response sensitivity. Lower index = more efficient.

Temp (°C)	SCAPE Index	vs 75°C ref
40°C	62.7×	0.92×
55°C	65.1×	0.95×
70°C	67.4×	0.99×
75°C ← now	**68.2×**	**1.00×**
85°C	69.8×	1.02×
95°C	71.3×	1.05×
100°C	72.1×	1.06×
105°C	72.8×	1.07×

EFFICIENCY SENS.

MODERATE

Apple M4 (2024) · n = 0.794 · nominal Tj = 75°C

Apple Silicon architecture has moderate temperature sensitivity at n=0.794. Halving operating temperature from 75°C improves efficiency by approximately 1.10×. The operating temperature range for Apple chips is narrower than x86 — M4 TDP of 28W means junction temperatures rarely approach TjMax under normal operation. The efficiency advantage over x86 holds across the full temperature range — Apple's structural gap cannot be closed by cooling.

EFFICIENCY FLOOR

LOW

Apple Silicon architectural efficiency floor · 4.55× above floor · floor = 15×

The structural efficiency floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this floor as a specific normalized SCAPE index.

The Apple Silicon architectural floor at ~15× is the information-theoretic minimum for ARM ISA switching overhead with Apple's unified memory integration at current nodes. At 68÷15 = 4.55× above floor, Apple has meaningful headroom within the current architecture. The floor is a system integration constraint, not a material or foundry constraint. What must happen next: system-level integration advances are the engineering lever. Node shrinks alone will not close the 4.5× gap to the architectural floor.

#3 Qualcomm Snapdragon X Elite

4nm · 2024 · 80W · 20.0B transistors · 3.8GHz

316×

Source: Qualcomm product page 2024. TSMC 4nm. ~20B transistors. 80W reference device TDP. 3.8 GHz Oryon cores.

HISTORY

2023	**Snapdragon 8cx Gen 3**	SCAPE 343× — Qualcomm ARM baseline on 4nm
2024	**Snapdragon X Elite**	SCAPE 316× — 7.9% improvement; ARM advantage over x86 confirmed

COMMENTARY

The X Elite at SCAPE 316× sits between Apple Silicon (65×) and x86 (570×). ARM ISA at the same TSMC 4nm node delivers 1.75× better efficiency than AMD. Apple's full system integration delivers a further 4.9× gap over Qualcomm. The X Elite ratio to AMD is 316÷576 = 0.55. The ARM structural floor is lower than x86 — each node step has more room to move at Snapdragon generation — ARM's lower structural floor gives each node step more room to move.

PREVIOUS GENERATION vs NOW

THEN Snapdragon 8cx Gen 3 (4nm, 2023)	SCAPE 343×	**NOW** SCAPE 316×	**-7.9%**

SCAPE read: The X Elite at SCAPE 316× sits between Apple Silicon (65×) and x86 (570×).

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
TDP	80 W	**PUBLISHED**
Transistor count	20.0 B	**PUBLISHED**
Base frequency	3.800 GHz	**PUBLISHED**
Process node	4nm	**PUBLISHED**
SCAPE Efficiency Index	315.9×	DERIVED: IAM
Distance above floor	9.03× (802.6% above)	DERIVED
Architecture n (temp. response)	0.500	DERIVED
Halving rate	12.7 nodes (if current step holds)	ILLUSTRATIVE
Most recent observed step	5.3%	OBSERVED
Illustrative SCAPE 2026	299×	ILLUSTRATIVE
Illustrative SCAPE 2028	283×	ILLUSTRATIVE
Illustrative SCAPE 2030	268×	ILLUSTRATIVE
Illustrative SCAPE 2032	254×	ILLUSTRATIVE
Distance to floor	40.4 nodes at current step rate	ILLUSTRATIVE
Thermal crossover	—	DERIVED
Global rank (7 chips)	#3 of 7	DERIVED

METRICS ANALYSIS

X Elite at SCAPE 316× ranks third globally. The ARM architecture advantage over x86 is structural and quantified: 1.82× at the same foundry, same node, same year. That gap cannot be closed by cooling, TDP optimization, or process refinement alone — it requires ISA-level redesign. Two things are analytically notable. First, the Qualcomm-to-Apple gap (316÷65 = 4.9×) is the system integration premium. Both are ARM. Both use TSMC. Qualcomm runs on Windows with a general-purpose OS stack. Apple runs on macOS with a fully co-designed software/hardware stack. That 4.9× gap is the measurable value of Apple's vertical integration — and it has not closed since M1 launched at TSMC 5nm in 2020. Second, Qualcomm's Oryon core is the first ARM-for-Windows architecture built ground-up for laptop efficiency rather than mobile power. The 8cx Gen 3 to X Elite improvement (343→316×, -7.9%) is meaningful at the same node. The 2nm Snapdragon generation is the first test of whether Oryon's architecture choices deliver a larger step than the single published result (5.3%) suggests. Illustrative 2028 scenario (current step held constant): 283×.

TRAJECTORY

Year	SCAPE Index	Rank	Milestone
2024	316×	3/7	
2026	299×	3/7	

Year	SCAPE Index	Rank	Milestone
2028	283×	4/7	
2030	268×	4/7	
2032	254×	4/7	

TEMPERATURE PREDICTIONS

SCAPE efficiency index across full junction temperature range. Derived from published specifications only — no internal data used. Architecture parameter n=0.500 determines temperature response sensitivity. Lower index = more efficient.

Temp (°C)	SCAPE Index	vs 75°C ref
40°C	299.6×	0.95×
55°C	306.7×	0.97×
70°C	313.6×	0.99×
75°C ← now	**315.9×**	**1.00×**
85°C	320.4×	1.01×
95°C	324.8×	1.03×
100°C	327.0×	1.04×
105°C	329.2×	1.04×

EFFICIENCY SENS.

LOW

Qualcomm X Elite (2024) · n = 0.500 · nominal Tj = 75°C

Qualcomm ARM architecture has low-to-moderate temperature sensitivity at n=0.500. Halving operating temperature improves efficiency by approximately 1.06×. The ARM efficiency advantage over x86 holds across the full operating temperature range — it is architectural, not thermal. At any junction temperature in the operating range, X Elite is more efficient than any x86 in this dataset.

EFFICIENCY FLOOR

MED

Qualcomm ARM structural efficiency floor · 9.03× above floor · floor = 35×

The structural efficiency floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this floor as a specific normalized SCAPE index.

The Qualcomm ARM architectural floor at ~35× is derived from the information-theoretic minimum for ARM ISA overhead without Apple's full system integration. At 316÷35 = 9.0× above the floor, substantial headroom remains. The gap between Qualcomm (316×) and Apple (65×) represents the system integration premium Apple's unified memory architecture and chip-package co-design delivers. What must happen next: deeper system integration is the engineering path to the floor. Qualcomm's Oryon core architecture moves in this direction. The 2nm Snapdragon generation tests whether ARM-for-Windows can close the gap to Apple.

#4 Intel Core Ultra 9 285K

3nm · 2024 · 125W · 17.8B transistors · 3.7GHz

570×

Source: Intel ARK product page; Tom's Hardware Oct 2024; TechPowerUp CPU Database. TSMC N3B Arrow Lake. 17.8B transistors. 125W TDP base. 3.7 GHz base clock.

HISTORY

2003	**Intel Pentium 4 3.2GHz**	SCAPE 151,776× — pre-Dennard; voltage scaling active
2006	**Intel Core 2 Duo E6600**	SCAPE 27,934× — Core 2 architecture reset; massive improvement
2011	**Intel Core i7-2600**	SCAPE 7,292× — Sandy Bridge; improvement continues but below Dennard rate
2015	**Intel Core i7-6700**	SCAPE 2,787× — Skylake 14nm; Intel 14nm plateau begins
2019	**Intel Core i9-9900KS**	SCAPE 2,376× — still on 14nm; five-year plateau
2021	**Intel Core i9-12900K**	SCAPE 2,228× — Alder Lake 10nm; finally off 14nm
2022	**Intel Core i9-13900K**	SCAPE 1,947× — Raptor Lake 10nm; power-hungry era
2024	**Intel Core Ultra 9 285K**	SCAPE 570× — Arrow Lake TSMC 3nm; largest Intel improvement in dataset

COMMENTARY

The 285K at SCAPE 570× is Intel's most significant efficiency improvement since Core 2 in 2006. The 70.7% improvement from 13900K is the largest single-generation improvement in the Intel dataset — and in the full 20-year x86 record. Intel's n=0.936 means it responds strongly to temperature. Intel leads AMD by 7.2% at 40°C cold. It trails AMD by 3.7% at 105°C TjMax. The crossover is 82°C. Below it, cooling investment returns value for Intel. Above it, the architecture works against you.

PREVIOUS GENERATION vs NOW

THEN Intel Core i9-13900K (Raptor Lake, 10nm, 2022)	SCAPE 1947×	**NOW** SCAPE 570×	**-70.7%**

SCAPE read: The 285K at SCAPE 570× is Intel's most significant efficiency improvement since Core 2 in 2006.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
TDP	125 W	**PUBLISHED**
Transistor count	17.8 B	**PUBLISHED**
Base frequency	3.700 GHz	**PUBLISHED**
Process node	3nm	**PUBLISHED**
SCAPE Efficiency Index	569.7×	DERIVED: IAM
Distance above floor	8.14× (713.9% above)	DERIVED
Architecture n (temp. response)	0.936	DERIVED
Halving rate	0.6 nodes (if current step holds)	ILLUSTRATIVE
Most recent observed step	70.7%	OBSERVED
Illustrative SCAPE 2026	167×	ILLUSTRATIVE
Illustrative SCAPE 2028	≤70× (floor)	ILLUSTRATIVE
Illustrative SCAPE 2030	≤70× (floor)	ILLUSTRATIVE
Illustrative SCAPE 2032	≤70× (floor)	ILLUSTRATIVE
Distance to floor	1.7 nodes at current step rate	ILLUSTRATIVE
Thermal crossover	~82°C (Intel leads below, AMD leads above)	DERIVED
Global rank (7 chips)	#4 of 7	DERIVED

METRICS ANALYSIS

The 285K at SCAPE 570× ranks fourth globally, just behind AMD at 576×. The Arrow Lake result deserves specific examination. The improvement from 1,947× to 570× is 70.7% — the largest single generational step in the Intel dataset across 20 years of published data. It breaks down roughly as: half from TSMC 3nm (node contribution), half from TDP discipline — 253W to 125W. Intel did not just change foundries. It changed its power philosophy. Two things are analytically notable. First, this step reversed a decade of near-flat improvement (5-15% per step) during the 14nm plateau. Whether this rate continues at 2nm is what the 2nm data point will reveal. Second, Intel's high n (0.936) is a double-edged property. Strong cooling

returns make Intel attractive for cold-running workloads. But sustained hot-load performance — AI inference, dense serving — runs above the 82°C crossover where AMD holds the advantage. Illustrative 2027 scenario (70.7% step held constant, one step): 167×. Note: projecting the 70.7% step two nodes gives 49× — below the x86 architectural floor of 70×. This is physically impossible. It is precisely why the floor signals this rate cannot be sustained.

TRAJECTORY

Year	SCAPE Index	Rank	Milestone
2024	570×	4/7	
2026	167×	2/7	
2028	70×	3/7	★ 2nm data: Dennard signal?
2030	70×	3/7	
2032	70×	3/7	

TEMPERATURE PREDICTIONS

SCAPE efficiency index across full junction temperature range. Derived from published specifications only — no internal data used. Architecture parameter n=0.936 determines temperature response sensitivity. Lower index = more efficient.

Temp (°C)	SCAPE Index	vs 75°C ref
40°C	515.9×	0.91×
55°C	539.0×	0.95×
70°C	562.0×	0.99×
75°C ← now	**569.7×**	**1.00×**
85°C	585.0×	1.03×
95°C	600.3×	1.05×
100°C	607.9×	1.07×
105°C	615.5×	1.08×

EFFICIENCY SENS.

HIGH

Intel 285K (2024) · n = 0.936 · nominal Tj = 75°C

Post-Dennard Intel architecture has high temperature sensitivity at n=0.936 — nearly linear response. Halving operating temperature improves efficiency by approximately 1.11×. Strong returns from aggressive cooling investment. But that sensitivity is a double edge: Intel degrades faster under sustained hot load. Below 82°C Intel leads AMD. Above 82°C AMD leads Intel. The crossover at 82°C junction temperature is a direct consequence of the n values applied to published specs — it is not a model assumption.

EFFICIENCY FLOOR

MED

x86 Post-Dennard Intel structural floor · 8.14× above floor · floor = 70×

The structural efficiency floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this floor as a specific normalized SCAPE index.

The post-Dennard x86 architectural floor at ~70× is the same for Intel and AMD — set by x86 ISA switching overhead, not by the foundry or the node. At 570÷70 = 8.1× above floor, Intel has substantial headroom. The 2nm data point is the next falsifiable test: The floor signals the per-step improvement will be materially below the 70.7% Arrow Lake result — the signal that the next Dennard-equivalent wall is registering. If the 2nm step matches Arrow Lake, the framework's regime-change signal is delayed. If it falls short, the framework's floor-derived signal is confirmed.

#5 AMD Ryzen 9 9950X

4nm · 2024 · 170W · 20.6B transistors · 4.3GHz

576×

Source: AMD official datasheet; Vortez Aug 2024; TechPowerUp CPU Database. 4nm TSMC Zen5 Granite Ridge. 17.2B CCD + 3.4B IOD = 20.6B total. 170W TDP. 4.3 GHz base.

HISTORY

2003	**AMD Athlon 64 3400+**	SCAPE 105,000× — pre-Dennard baseline; full voltage scaling era
2007	**AMD Phenom X4 9600**	SCAPE 27,951× — Dennard breakdown visible; improvement rate slows
2011	**AMD FX-8150 Bulldozer**	SCAPE 8,760× — Bulldozer era; efficiency improving but below Dennard projection
2017	**AMD Ryzen 7 1800X**	SCAPE 1,650× — Zen architecture resets AMD efficiency baseline
2019	**AMD Ryzen 9 3900X**	SCAPE 839× — Zen2 7nm; AMD holds 2.5× efficiency lead over Intel 9900KS
2020	**AMD Ryzen 9 5950X**	SCAPE 946× — Zen3 7nm; slight regression (higher TDP, lower freq)
2022	**AMD Ryzen 9 7950X**	SCAPE 863× — Zen4 5nm; improvement resumes
2024	**AMD Ryzen 9 9950X**	SCAPE 576× — Zen5 4nm; AMD lead closes to <1% vs Intel

COMMENTARY

AMD at 576× holds the efficiency lead over Intel at nominal temperature by less than 1%. The 2.5× efficiency advantage AMD held over Intel's best in 2019 has compressed to statistical noise. Two things are worth noting. First, AMD's n=0.340 means it degrades slowly with heat — above 82°C AMD leads Intel, and sustained AI inference runs above 80°C. The workload that matters most for hyperscalers is the one AMD already wins. Second, AMD and Apple share the same TSMC foundry at the same node. AMD is 8.5× less efficient. That gap is the entire x86 ISA overhead.

PREVIOUS GENERATION vs NOW

THEN AMD Ryzen 9 7950X (Zen4, 5nm, 2022)	SCAPE 863×	**NOW** SCAPE 576×	**-33.3%**

SCAPE read: AMD at 576× holds the efficiency lead over Intel at nominal temperature by less than 1%.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
TDP	170 W	**PUBLISHED**
Transistor count	20.6 B	**PUBLISHED**
Base frequency	4.300 GHz	**PUBLISHED**
Process node	4nm	**PUBLISHED**
SCAPE Efficiency Index	576.0×	DERIVED: IAM
Distance above floor	8.23× (722.9% above)	DERIVED
Architecture n (temp. response)	0.340	DERIVED
Halving rate	1.7 nodes (if current step holds)	ILLUSTRATIVE
Most recent observed step	33.2%	OBSERVED
Illustrative SCAPE 2026	385×	ILLUSTRATIVE
Illustrative SCAPE 2028	257×	ILLUSTRATIVE
Illustrative SCAPE 2030	172×	ILLUSTRATIVE
Illustrative SCAPE 2032	115×	ILLUSTRATIVE
Distance to floor	5.2 nodes at current step rate	ILLUSTRATIVE
Thermal crossover	~82°C (AMD leads above, Intel leads below)	DERIVED
Global rank (7 chips)	#5 of 7	DERIVED

METRICS ANALYSIS

The 9950X at SCAPE 576× ranks fifth globally and first among x86. The gap to Intel's 285K is 576 vs 570 — less than 1.1%. Five years of AMD efficiency leadership built through Zen2 and Zen3 has compressed to statistical noise. Two things stand out analytically. First, the AMD-Intel convergence at 4nm is not accidental. Intel moved to TSMC 3nm and cut TDP from 253W to 125W — a foundry change and a TDP discipline decision simultaneously. AMD held at 4nm Zen5 at 170W. The convergence is a story about Intel's choices, not AMD's decline. Second, AMD's

temperature advantage is structural and workload-relevant. At 75°C AMD trails Intel by 1%. At 105°C TjMax, AMD leads by 3.7%. The crossover is at 82°C. Most data center AI inference workloads run at or above 80°C under load. AMD already wins the sustained-load workload that drives hyperscaler procurement decisions. No cooling investment changes that — the advantage comes from AMD's lower n, not from temperature. Illustrative 2028 scenario (current step held constant): 385×. The 2nm data point will show whether Intel's larger recent step translates to a widening gap.

TRAJECTORY

Year	SCAPE Index	Rank	Milestone
2024	576×	5/7	
2026	385×	5/7	
2028	257×	3/7	★ 2nm data: Dennard signal?
2030	172×	4/7	
2032	115×	4/7	

TEMPERATURE PREDICTIONS

SCAPE efficiency index across full junction temperature range. Derived from published specifications only — no internal data used. Architecture parameter n=0.340 determines temperature response sensitivity. Lower index = more efficient.

Temp (°C)	SCAPE Index	vs 75°C ref
40°C	555.6×	0.96×
55°C	564.5×	0.98×
70°C	573.2×	1.00×
75°C ← now	**576.0×**	**1.00×**
85°C	581.6×	1.01×
95°C	587.0×	1.02×
100°C	589.7×	1.02×
105°C	592.4×	1.03×

EFFICIENCY SENS.

MODERATE

AMD 9950X (2024) · n = 0.340 · nominal Tj = 75°C

Post-Dennard AMD architecture responds moderately to junction temperature at n=0.340. Halving operating temperature from 75°C improves efficiency by approximately 1.04×. The AMD advantage is not from cooling sensitivity — it is from slower hot-side degradation. Above 82°C AMD is more efficient than Intel. Below 82°C Intel leads. For sustained AI inference workloads running at 80-105°C, AMD holds the efficiency advantage.

EFFICIENCY FLOOR

MED

x86 Post-Dennard AMD structural floor · 8.23× above floor · floor = 70×

The structural efficiency floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this floor as a specific normalized SCAPE index.

The post-Dennard x86 architectural floor at ~70× is derived from the information-theoretic minimum for the x86 ISA switching overhead at current CMOS node geometries. At 576÷70 = 8.2× above floor, AMD has substantial headroom within the current architecture. The floor is not a material constraint — it is an ISA-level architectural constraint. Closing the gap to Apple Silicon (68×) requires ISA redesign, not a foundry change. AMD and Apple share the same TSMC foundry. The 8.5× gap is pure architecture. The next published AMD data point shows whether the floor-derived signal is registering.

#6 NVIDIA H200 SXM

4nm · 2024 · 700W · 80.0B transistors · 1.98GHz

1326×

Source: NVIDIA H200 SXM product page, 2024. TSMC N4. 80.0B transistors — identical GH100 die as H100 (confirmed: developer.nvidia.com/blog/nvidia-hopper-architecture-in-depth). 700W TDP SXM. SM boost clock 1.98 GHz (topcpu.net / videocardz.net).

HISTORY

2022	**NVIDIA H100 SXM5**	SCAPE 1,435× — Hopper architecture; 4nm baseline
2024	**NVIDIA H200 SXM**	SCAPE 1,326× — Hopper variant; 7.6% improvement from H100

COMMENTARY

H200 at SCAPE 1,326× shows 7.6% improvement over H100 at the same 4nm node. No node shrink and identical compute die — GH100 silicon is unchanged. The H200 upgrade is entirely in the memory subsystem: 141GB HBM3e at 4.8 TB/s versus H100's 80GB HBM3 at 3.35 TB/s. The SCAPE framework correctly reflects this: compute switching efficiency improves only modestly (higher SM boost clock 1.98 vs 1.83 GHz), while the memory advantage is invisible to SCAPE — which measures switching energy, not bandwidth.

PREVIOUS GENERATION vs NOW

THEN NVIDIA H100 SXM5 (4nm, 2022)	SCAPE 1435×	**NOW** SCAPE 1326×	**-7.6%**

SCAPE read: H200 at SCAPE 1,326× shows 7.6% improvement over H100 at the same 4nm node.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
TDP	700 W	**PUBLISHED**
Transistor count	80.0 B	**PUBLISHED**
Base frequency	1.980 GHz	**PUBLISHED**
Process node	4nm	**PUBLISHED**
SCAPE Efficiency Index	1326.4×	DERIVED: IAM
Distance above floor	5.31× (430.6% above)	DERIVED
Architecture n (temp. response)	0.650	DERIVED
Halving rate	8.2 nodes (if current step holds)	ILLUSTRATIVE
Most recent observed step	8.1%	OBSERVED
Illustrative SCAPE 2026	1219×	ILLUSTRATIVE
Illustrative SCAPE 2028	1120×	ILLUSTRATIVE
Illustrative SCAPE 2030	1030×	ILLUSTRATIVE
Illustrative SCAPE 2032	946×	ILLUSTRATIVE
Distance to floor	19.8 nodes at current step rate	ILLUSTRATIVE
Thermal crossover	—	DERIVED
Global rank (7 chips)	#6 of 7	DERIVED

METRICS ANALYSIS

H200 at SCAPE 1,326× sits just above H100 at 1,435×, with a 7.6% compute-efficiency gain. Two things are analytically important here. First, the small SCAPE delta between H100 and H200 is not a weakness — it is the framework correctly identifying that H200 is the same GH100 die with a larger memory stack. The H200's real-world advantage over H100 (20-40% faster inference on large models) comes entirely from memory bandwidth, not switching efficiency. SCAPE measures switching efficiency. A bandwidth-focused upgrade should produce a small SCAPE delta — and it does. Second, the NVIDIA GPU trajectory gets its next genuine data point with the Blackwell successor — a new compute die rather than a memory upgrade. The Blackwell successor represents the first genuine new die since Hopper. If that step delivers a SCAPE improvement larger than AMD or Intel achieve at 2nm, GPU architecture is outpacing CPU on the normalized scale. The number to watch: SCAPE improvement from H200 (1,326×) to next NVIDIA flagship. Illustrative 2028 scenario (current step held constant): 1120×.

TRAJECTORY

Year	SCAPE Index	Rank	Milestone
2024	1326×	6/7	

Year	SCAPE Index	Rank	Milestone
2026	1219×	6/7	
2028	1120×	6/7	★ Blackwell successor
2030	1030×	6/7	

TEMPERATURE PREDICTIONS

SCAPE efficiency index across full junction temperature range. Derived from published specifications only — no internal data used. Architecture parameter n=0.650 determines temperature response sensitivity. Lower index = more efficient.

Temp (°C)	SCAPE Index	vs 75°C ref
40°C	1238.1×	0.93×
55°C	1276.4×	0.96×
70°C	1314.0×	0.99×
75°C ← now	**1326.4×**	**1.00×**
85°C	1351.0×	1.02×
95°C	1375.4×	1.04×
100°C	1387.6×	1.05×
105°C	1399.6×	1.06×

EFFICIENCY SENS.

MODERATE

NVIDIA H200 SXM (2024) · n = 0.650 · nominal Tj = 75°C

Same NVIDIA GPU temperature response profile as H100: n=0.650. Halving operating temperature improves efficiency by approximately 1.08×. GPU temperature management is a system-level engineering challenge — the 700W TDP ceiling constrains how much cooling can do at the chip level. Rack-level liquid cooling changes the effective operating point.

EFFICIENCY FLOOR

MED

NVIDIA GPU throughput architecture floor · 5.31× above floor · floor = 250×

The structural efficiency floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this floor as a specific normalized SCAPE index.

The GPU throughput architectural floor at ~250× is derived from the information-theoretic minimum for sustained parallel compute at this transistor density. H200 at 1,326÷250 = 5.31× above the floor has meaningful headroom remaining. The modest H100-to-H200 SCAPE improvement (7.6%) reflects a memory upgrade, not a compute architecture advance. The next genuine compute step is Blackwell — that is when the SCAPE trajectory for NVIDIA resumes meaningful movement. The Blackwell successor data point will show whether GPU compute efficiency is accelerating on the normalized scale — the first genuine new compute die since Hopper.

#7 NVIDIA H100 SXM5

4nm · 2022 · 700W · 80.0B transistors · 1.83GHz

1435×

Source: NVIDIA Hopper Architecture In-Depth, developer.nvidia.com, March 2022. TSMC N4. 80B transistors GH100. 700W TDP SXM5. SM boost clock 1.83 GHz.

HISTORY

2016	**NVIDIA GTX 1080 (Pascal)**	SCAPE 8,214× — GPU baseline; throughput architecture
2018	**NVIDIA RTX 2080 (Turing)**	SCAPE 3,151× — first RT cores; efficiency improving
2020	**NVIDIA A100 (Ampere)**	SCAPE 1,527× — data center GPU class; 7nm
2022	**NVIDIA H100 SXM5 (Hopper)**	SCAPE 1,435× — 4nm; 6.0% improvement from A100

COMMENTARY

H100 at SCAPE 1,435× ranks last on the efficiency index — but the comparison requires context. GPU architecture trades per-transistor switching efficiency for massive parallelism. H100 runs 80 billion transistors at 700W sustaining thousands of simultaneous matrix operations. The SCAPE index measures per-transistor switching overhead, not throughput. The relevant question is not where H100 sits on the CPU efficiency scale. It is whether GPU efficiency is improving per generation — and at what rate.

PREVIOUS GENERATION vs NOW

THEN NVIDIA A100 (7nm, 2020)	SCAPE 1527×	**NOW** SCAPE 1435×	**-6.0%**

SCAPE read: H100 at SCAPE 1,435× ranks last on the efficiency index — but the comparison requires context.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
TDP	700 W	**PUBLISHED**
Transistor count	80.0 B	**PUBLISHED**
Base frequency	1.830 GHz	**PUBLISHED**
Process node	4nm	**PUBLISHED**
SCAPE Efficiency Index	1435.1×	DERIVED: IAM
Distance above floor	5.74× (474.0% above)	DERIVED
Architecture n (temp. response)	0.650	DERIVED
Halving rate	8.2 nodes (if current step holds)	ILLUSTRATIVE
Most recent observed step	8.1%	OBSERVED
Illustrative SCAPE 2026	1319×	ILLUSTRATIVE
Illustrative SCAPE 2028	1212×	ILLUSTRATIVE
Illustrative SCAPE 2030	1114×	ILLUSTRATIVE
Illustrative SCAPE 2032	1024×	ILLUSTRATIVE
Distance to floor	20.7 nodes at current step rate	ILLUSTRATIVE
Thermal crossover	—	DERIVED
Global rank (7 chips)	#7 of 7	DERIVED

METRICS ANALYSIS

H100 at SCAPE 1,435× represents the Hopper architecture baseline. The 6.0% improvement from A100 (1,527→1,435×) reflects a generation that prioritized compute throughput over switching efficiency — a rational engineering trade-off for AI workloads. Two things are analytically notable. First, the H200 result clarifies the H100 story: H200 shows only 7.6% improvement because it is the same GH100 die — the H200 upgrade is memory bandwidth, not compute. The Blackwell successor (new die) is when NVIDIA's SCAPE trajectory resumes real movement. SCAPE improvement than AMD or Intel's 2nm step. If NVIDIA's improvement at Blackwell exceeds the AMD/Intel 2nm improvement the Blackwell successor data point will show whether GPU improves faster than CPU on the floor-derived scale. Illustrative 2028 scenario (current step held constant): 1212×.

TRAJECTORY

Year	SCAPE Index	Rank	Milestone
2022	1435×	7/7	

Year	SCAPE Index	Rank	Milestone
2024	1435×	7/7	
2026	1319×	7/7	
2028	1212×	7/7	
2030	1114×	7/7	

TEMPERATURE PREDICTIONS

SCAPE efficiency index across full junction temperature range. Derived from published specifications only — no internal data used. Architecture parameter n=0.650 determines temperature response sensitivity. Lower index = more efficient.

Temp (°C)	SCAPE Index	vs 75°C ref
40°C	1339.6×	0.93×
55°C	1381.0×	0.96×
70°C	1421.7×	0.99×
75°C ← now	**1435.1×**	**1.00×**
85°C	1461.8×	1.02×
95°C	1488.2×	1.04×
100°C	1501.3×	1.05×
105°C	1514.3×	1.06×

EFFICIENCY SENS.

MODERATE

NVIDIA H100 SXM5 (2022) · n = 0.650 · nominal Tj = 75°C

NVIDIA GPU architecture has moderate temperature sensitivity at n=0.650. Halving operating temperature improves efficiency by approximately 1.08×. GPU junction temperatures under sustained compute load routinely reach 80-83°C. The SCAPE index is per-transistor switching efficiency — GPU efficiency for specific workloads should be evaluated as ops-per-watt, not SCAPE index alone. The framework applies to GPU architecture as a normalized cross-platform benchmark.

EFFICIENCY FLOOR

MED

NVIDIA GPU throughput architecture floor · 5.74× above floor · floor = 250×

The structural efficiency floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this floor as a specific normalized SCAPE index.

The GPU throughput architectural floor at ~250× reflects the information-theoretic minimum for driving thousands of parallel compute units at sustained high load. At 1,435÷250 = 5.74× above this floor, H100 has substantial headroom. The improvement path for GPU architecture runs through higher transistor utilization per watt — Blackwell addresses this directly with disaggregated compute and memory. What must happen next: the Blackwell successor data point shows whether GPU compute efficiency is accelerating or plateauing on the normalized scale.

THE COMPLETE 20-YEAR DATASET — ALL DERIVED OUTPUTS

Three published inputs per chip — TDP, transistors, frequency — produce every value below. Orange = published inputs. Teal = derived. Sources cited by superscript in the reference table on the following page.

TABLE 1 — PUBLISHED INPUTS AND CORE DERIVED VALUES

Orange columns = three published inputs per chip. Teal columns = values derived from those three inputs by the SCAPE framework. No proprietary data used.

Chip	Yr	Node	TDP (W) PUBL	Trans (B) PUBL	Freq (GHz) PUBL	SCAPE Index @75°C	Archn	Ref
AMD Athlon 64 3400+	2003	130nm	89	0.106	2.4	105002×	0.340	[1]
AMD Athlon 64 X2 3800+	2005	90nm	89	0.154	2.0	86729×	0.340	[2]
AMD Phenom X4 9600	2007	65nm	95	0.45	2.3	27549×	0.340	[3]
AMD Phenom II X4 955	2009	45nm	125	0.758	3.2	15467×	0.340	[4]
AMD FX-8150 Bulldozer	2011	32nm	125	1.2	3.6	8685×	0.340	[5]
AMD Ryzen 7 1800X	2017	14nm	95	4.8	3.6	1650×	0.340	[6]
AMD Ryzen 9 3900X	2019	7nm	105	9.89	3.8	839×	0.340	[7]
AMD Ryzen 9 5950X	2020	7nm	105	9.8	3.4	946×	0.340	[8]
AMD Ryzen 9 7950X	2022	5nm	170	13.14	4.5	863×	0.340	[9]
AMD Ryzen 9 9950X	2024	4nm	170	20.6	4.3	576×	0.340	[10]
Intel Pentium 4 3.2GHz	2003	130nm	89	0.055	3.2	151776×	0.936	[11]
Intel Core 2 Duo E6600	2006	65nm	65	0.291	2.4	27934×	0.936	[12]
Intel Core i7-920	2008	45nm	130	0.731	2.67	19991×	0.936	[13]
Intel Core i7-2600	2011	32nm	95	1.16	3.4	7230×	0.936	[14]
Intel Core i7-3770	2012	22nm	77	1.4	3.4	4855×	0.936	[15]
Intel Core i7-6700	2015	14nm	65	1.75	4.0	2787×	0.936	[16]
Intel Core i9-9900KS	2019	14nm	95	3.0	4.0	2376×	0.936	[17]
Intel Core i9-12900K	2021	10nm	241	10.0	3.2	2260×	0.936	[18]
Intel Core i9-13900K	2022	10nm	253	13.0	3.0	1947×	0.936	[19]
Intel Core Ultra 9 285K	2024	3nm	125	17.8	3.7	570×	0.936	[20]
Apple M1	2020	5nm	15	16.0	3.2	88×	0.794	[21]
Apple M2	2022	5nm	15	20.0	3.49	64×	0.794	[22]
Apple M3	2023	3nm	22	25.0	4.05	65×	0.794	[23]
Apple M4	2024	3nm	28	28.0	4.4	68×	0.794	[24]
Snapdragon 8cx Gen 3	2023	4nm	45	13.5	3.0	334×	0.500	[25]
Snapdragon X Elite	2024	4nm	80	20.0	3.8	316×	0.500	[26]
NVIDIA A100 (GA100)	2020	7nm	400	54.2	1.41	1571×	0.650	[27]
NVIDIA H100 SXM5	2022	4nm	700	80.0	1.83	1435×	0.650	[28]
NVIDIA H200 SXM	2024	4nm	700	80.0	1.98	1326×	0.650	[29]

■ Orange = published input ■ Teal = IAMPerformance derived output * Three published numbers per chip produce every derived value in this table.

THE COMPLETE 20-YEAR DATASET — THERMAL AND DIRECTIONAL ANALYSIS

Every thermal output derived from three published inputs and architecture parameter n. SCAPE at cold (40°C) and hot (105°C) are physics-derived measurements, not projections. Illustrative scenarios use most recent observed improvement step — not a validated forward rate. x86 Gap = how many times less efficient than Apple M3 (65×) — x86 reference only, Apple rows show dashes.

Chip	Yr	SCAPE @40°C (cold)	SCAPE @75°C (nominal)	SCAPE @105°C (TjMax)	Arch n	Recent Step	Illustrative 2028	Illustrative 2032	Floor Ratio	x86 Gap vs M3 (65×)	Ref
AMD Athlon 64 3400+	2003	101287×	105002×	107995×	0.340	33%/nd	46854×	20908×	1500.03×	1610.5×	[1]
AMD Athlon 64 X2 3800+	2005	83660×	86729×	89201×	0.340	33%/nd	38701×	17269×	1238.99×	1330.2×	[2]
AMD Phenom X4 9600	2007	26574×	27549×	28334×	0.340	33%/nd	12293×	5485×	393.56×	422.5×	[3]
AMD Phenom II X4 955	2009	14920×	15467×	15908×	0.340	33%/nd	6902×	3080×	220.96×	237.2×	[4]
AMD FX-8150 Bulldozer	2011	8377×	8685×	8932×	0.340	33%/nd	3875×	1729×	124.07×	133.2×	[5]
AMD Ryzen 7 1800X	2017	1592×	1650×	1697×	0.340	33%/nd	736×	329×	23.57×	25.3×	[6]
AMD Ryzen 9 3900X	2019	809×	839×	862×	0.340	33%/nd	374×	167×	11.98×	12.9×	[7]
AMD Ryzen 9 5950X	2020	912×	946×	973×	0.340	33%/nd	422×	188×	13.51×	14.5×	[8]
AMD Ryzen 9 7950X	2022	832×	863×	888×	0.340	33%/nd	385×	172×	12.33×	13.2×	[9]
AMD Ryzen 9 9950X	2024	556×	576×	592×	0.340	33%/nd	257×	115×	8.23×	8.8×	[10]
Intel Pentium 4 3.2GHz	2003	137446×	151776×	163984×	0.936	71%/nd	13030×	1119×	2168.22×	2327.8×	[11]
Intel Core 2 Duo E6600	2006	25297×	27934×	30181×	0.936	71%/nd	2398×	206×	399.06×	428.4×	[12]
Intel Core i7-920	2008	18104×	19991×	21599×	0.936	71%/nd	1716×	147×	285.59×	306.6×	[13]
Intel Core i7-2600	2011	6547×	7230×	7811×	0.936	71%/nd	621×	70×	103.28×	110.9×	[14]
Intel Core i7-3770	2012	4397×	4855×	5246×	0.936	71%/nd	417×	70×	69.36×	74.5×	[15]
Intel Core i7-6700	2015	2524×	2787×	3011×	0.936	71%/nd	239×	70×	39.81×	42.7×	[16]
Intel Core i9-9900KS	2019	2152×	2376×	2567×	0.936	71%/nd	204×	70×	33.94×	36.4×	[17]
Intel Core i9-12900K	2021	2047×	2260×	2442×	0.936	71%/nd	194×	70×	32.29×	34.7×	[18]
Intel Core i9-13900K	2022	1763×	1947×	2104×	0.936	71%/nd	167×	70×	27.82×	29.9×	[19]
Intel Core Ultra 9 285K	2024	516×	570×	616×	0.936	71%/nd	70×	70×	8.14×	8.7×	[20]
Apple M1	2020	81×	88×	94×	0.794	27%/nd	47×	26×	5.86×	—	[21]
Apple M2	2022	59×	64×	69×	0.794	27%/nd	35×	19×	4.30×	—	[22]
Apple M3	2023	60×	65×	70×	0.794	27%/nd	35×	19×	4.35×	—	[23]
Apple M4	2024	63×	68×	73×	0.794	27%/nd	37×	20×	4.55×	—	[24]
Snapdragon 8cx Gen 3	2023	316×	334×	348×	0.500	5%/nd	299×	268×	9.53×	5.1×	[25]
Snapdragon X Elite	2024	300×	316×	329×	0.500	5%/nd	283×	254×	9.03×	4.8×	[26]
NVIDIA A100 (GA100)	2020	1466×	1571×	1658×	0.650	8%/nd	1327×	1121×	6.28×	24.1×	[27]
NVIDIA H100 SXM5	2022	1340×	1435×	1514×	0.650	8%/nd	1212×	1024×	5.74×	22.0×	[28]
NVIDIA H200 SXM	2024	1238×	1326×	1400×	0.650	8%/nd	1120×	946×	5.31×	20.3×	[29]

All teal values derived from published inputs only — no proprietary data used. SCAPE@40°C and SCAPE@105°C are physics-derived measurements from architecture parameter n — not projections. Floor Ratio = current SCAPE ÷ structural efficiency floor — how far above the architectural floor. Illustrative 2028/2032 = constant-step scenario using most recent observed improvement step — not a forward prediction. x86 Gap vs M3 = SCAPE ÷ 65 — how many times less efficient than Apple M3 at nominal temperature. Apple rows show — (not applicable). Methodology protected under Patent Applications 64/012,720 and 64/014,568.

DATA SOURCES — 20-YEAR DATASET

Every published specification used in this analysis. Superscript reference numbers correspond to the [Ref] column in the dataset tables above.

[#]	Author / Publisher	Title	URL / Date	Specifications Used
[1]	Advanced Micro Devices	**AMD Athlon 64 3400+ Processor Product Page**	amd.com, 2004	130nm, 105.9M transistors, 89W TDP, 2.4 GHz, Socket 939
[2]	Advanced Micro Devices	**AMD Athlon 64 X2 3800+ Processor Specifications**	techpowerup.com/cpu-specs, 2005	90nm Manchester die, 154M transistors, 89W TDP, 2.0 GHz, Socket 939
[3]	Advanced Micro Devices	**AMD Phenom X4 9600 Processor Specifications**	techpowerup.com/cpu-specs, 2007	65nm Agena die, 450M transistors, 95W TDP, 2.3 GHz, Socket AM2+
[4]	Scott Wasson	**AMD Phenom II X4 955 Black Edition Review**	tomshardware.com/reviews/phenom-x4-955,2278.html, April 23, 2009	45nm Deneb die, 758M transistors, 125W TDP, 3.2 GHz
[5]	Ryan Smith	**AMD FX-8150 Bulldozer Review**	anandtech.com/show/4955, October 2011	32nm, 1.2B transistors (corrected by AMD PR), 125W TDP, 3.6 GHz
[6]	Advanced Micro Devices	**AMD Ryzen 7 1800X Product Page**	amd.com/en/products/cpu/amd-ryzen-7-1800x, 2017	14nm GlobalFoundries, 4.8B transistors Zeppelin die, 95W TDP, 3.6 GHz, Socket AM4
[7]	Ian Cutress	**AMD Ryzen 9 3900X Review**	anandtech.com/show/14605, July 2019	7nm TSMC, 9.89B transistors total (2× CCD + IOD), 105W TDP, 3.8 GHz
[8]	Advanced Micro Devices	**AMD Ryzen 9 5950X Product Specifications**	amd.com/en/products/cpu/amd-ryzen-9-5950x, 2020	7nm TSMC Zen3 Vermeer, 9.8B transistors, 105W TDP, 3.4 GHz base
[9]	Vortez	**AMD Ryzen 9 7950X Review**	vortez.net/articles_pages/amd_ryzen_9_7950x_review, 2022	5nm TSMC Zen4 Raphael, 13.14B transistors total, 170W TDP, 4.5 GHz
[10]	Vortez	**AMD Ryzen 9 9950X and 9900X Review**	vortez.net/articles_pages/amd_ryzen_9_9950x_amp_9900x_review, 2024	4nm TSMC Zen5, 17.2B CCD + 3.4B IOD = 20.6B total, 170W TDP, 4.3 GHz
[11]	Intel Corporation	**Intel Pentium 4 3.20 GHz Product Specifications**	ark.intel.com, 2003	130nm Northwood C1, 55M transistors, 89W TDP, 3.2 GHz
[12]	Intel Corporation	**Intel Core 2 Duo E6600 Product Specifications**	ark.intel.com, 2006	65nm Conroe, 291M transistors, 65W TDP, 2.4 GHz
[13]	Intel Corporation	**Intel Core i7-920 Processor Product Specifications**	ark.intel.com, 2008	45nm Bloomfield, 731M transistors, 130W TDP, 2.67 GHz
[14]	Intel Corporation	**Intel Core i7-2600 Processor Product Specifications**	ark.intel.com, 2011	32nm Sandy Bridge, 1.16B transistors, 95W TDP, 3.4 GHz
[15]	Intel Corporation	**Intel Core i7-3770 Processor Product Specifications**	ark.intel.com, 2012	22nm Ivy Bridge, 1.4B transistors, 77W TDP, 3.4 GHz
[16]	Intel Corporation	**Intel Core i7-6700 Processor Product Specifications**	ark.intel.com, 2015	14nm Skylake, 1.75B transistors, 65W TDP, 4.0 GHz
[17]	Intel Corporation	**Intel Core i9-9900KS Processor Product Specifications**	ark.intel.com, 2019	14nm Coffee Lake-S, 3.0B transistors, 95W TDP, 4.0 GHz base
[18]	Intel Corporation	**Intel Core i9-12900K Processor Product Specifications**	ark.intel.com, 2021	Intel 7 (10nm), approx. 10B transistors, 241W TDP, 3.2 GHz base
[19]	TechPowerUp CPU Database	**Intel Core i9-13900K Specifications**	techpowerup.com/cpu-specs, 2024	Intel 7 (10nm) Raptor Lake, 13.0B transistors, 253W TDP, 3.0 GHz base
[20]	Paul Alcorn	**Intel Core Ultra 9 285K and Core Ultra 5 245K Review**	tomshardware.com/pc-components/cpus/intel-core-ultra-9-285k-cpu-review, October 24, 2024	TSMC N3B, 17.8B transistors, 125W TDP
[21]	Apple Inc.	**Apple Unleashes M1**	apple.com/newsroom/2020/11/apple-unleashes-m1, November 2020	TSMC 5nm, 16B transistors, chip TDP approx. 15W, 3.2 GHz P-core
[22]	Apple Inc.	**Apple Unveils M2 Chip**	apple.com/newsroom, June 2022	TSMC N5P, 20B transistors, 15W chip TDP, 3.49 GHz P-core
[23]	Apple Inc.	**Apple Unveils M3, M3 Pro, and M3 Max**	apple.com/newsroom/2023/10/apple-unveils-m3-m3-pro-and-m3-max, October 2023	TSMC N3B, M3: 25B transistors, 22W chip TDP, 4.05 GHz P-core
[24]	Apple Inc.	**Apple Introduces M4 Chip**	apple.com/newsroom/2024/05/apple-introduces-m4-chip, May 2024	TSMC N3E, 28B transistors, 28W TDP, 4.4 GHz P-core
[25]	Qualcomm Technologies	**Snapdragon 8cx Gen 3 Compute Platform Product Brief**	qualcomm.com, 2023	TSMC 4nm, approx. 13.5B transistors, 45W TDP, 3.0 GHz Oryon cores
[26]	Qualcomm Technologies	**Snapdragon X Elite Platform Product Brief**	qualcomm.com/products/mobile/snapdragon/laptops-and-tablets/snapdragon-x-series/snapdragon-x-elite, 2024	TSMC 4nm, approx. 20B transistors, 80W reference device TDP, 3.8 GHz Oryon cores
[27]	NVIDIA Corporation	**NVIDIA A100 Tensor Core GPU Architecture Whitepaper**	images.nvidia.com/aem-dam/.../nvidia-ampere-architecture-whitepaper.pdf, 2020	TSMC 7nm, 54.2B transistors GA100, 400W TDP, SM boost clock 1.41 GHz
[28]	NVIDIA Corporation	**NVIDIA Hopper Architecture In-Depth**	developer.nvidia.com/blog/nvidia-hopper-architecture-in-depth, March 2022	TSMC N4, 80B transistors GH100, 700W TDP SXM5, SM boost clock 1.83 GHz
[29]	NVIDIA Corporation	**NVIDIA H200 Tensor Core GPU Product Specifications**	nvidia.com/en-us/data-center/h200, 2024	TSMC N4, 80.0B transistors (identical GH100 die as H100), 700W TDP SXM, SM boost clock 1.98 GHz

All specifications are from primary published sources. No proprietary or internal data was used or accessed in producing this analysis. Methodology protected under Patent Applications 64/012,720 and 64/014,568.

CONCLUSION

Every chip ever designed has a physical floor — a limit set not by engineering choices but by the thermodynamic cost of information itself. That floor is the same for every architecture. It does not care whether the transistor is x86, ARM, GPU, or anything that comes after. It is set by temperature, and nothing else. IAM's Law derives it from first principles. The SCAPE index measures how far above it each chip operates. For the first time, AMD at 576×, Intel at 570×, Apple Silicon at 65×, Qualcomm at 316×, and NVIDIA at 1,326× appear on one scale — not because someone decided to rank them, but because physics gives them no choice but to share it.

The floor is not just a ranking tool. It is a detection instrument. When the most recent published improvement step would project an architecture through its own physical floor within a small number of nodes, that trajectory is physically impossible. An impossible trajectory means the rate cannot be sustained. That is the regime change signal — the same signal the framework detected from two data points in 2003-2005, before the Dennard breakdown was widely acknowledged by the industry. Intel's Arrow Lake at 70.7% improvement projects the x86 floor reached in 1.7 nodes. The floor cannot be crossed. The same instrument that called the last transition is calling the next one. The 2nm data point is the live test.

The semiconductor industry understands its engineering limits with great sophistication. What has not existed until now is a fixed physical reference point — derived from first principles, not fitted to data — that makes both limits visible simultaneously: where each architecture stands relative to the thermodynamic floor, and when the current engineering path is running out of room. Those are two different questions. The SCAPE framework answers both from three published numbers. No benchmarks. No vendor data. No internal access required.

Issue 001 is the baseline. Every subsequent issue adds a data point to the prediction record. The framework's credibility is not claimed — it is accumulated. When 2nm data publishes, the floor-derived signal is either confirmed or not. When the Blackwell successor publishes, GPU trajectory gets its next data point. When M5 publishes, the Apple architectural gap gets its next measurement. The record builds in public, against the dated floor-derived signal, from primary sources. That is what a physical instrument looks like over time.

REFERENCES

[1] Advanced Micro Devices. AMD Ryzen 9 9950X Product Specifications. amd.com, 2024. 4nm TSMC Zen5, 20.6B transistors, 170W TDP.
[2] Alcorn, P. Intel Core Ultra 9 285K Review. Tom's Hardware, October 24, 2024. TSMC N3B, 17.8B transistors, 125W TDP.
[3] Apple Inc. Apple Introduces M4 Chip. apple.com/newsroom, May 2024. TSMC N3E, 28B transistors, 28W TDP, 4.4 GHz.
[4] Apple Inc. Apple Unveils M3. apple.com/newsroom, October 2023. TSMC N3B, 25B transistors, 22W TDP, 4.05 GHz.
[5] Apple Inc. Apple Unleashes M1. apple.com/newsroom, November 2020. TSMC 5nm, 16B transistors, 15W TDP, 3.2 GHz.
[6] Cutress, I. AMD Ryzen 9 3900X Review. AnandTech, July 2019. 7nm TSMC, 9.89B transistors, 105W TDP, 3.8 GHz.
[7] Dennard, R.H. et al. Design of ion-implanted MOSFET's with very small physical dimensions. IEEE JSSC, 1974.
[8] Intel Corporation. Intel Core Ultra 9 Processor 285K Specifications. ark.intel.com. TSMC N3B, 17.8B transistors, 125W.
[9] Intel Corporation. Intel Core i9 processor family specifications (i7-920 through i9-13900K). ark.intel.com, 2008-2022.
[10] NVIDIA Corporation. NVIDIA Hopper Architecture In-Depth. developer.nvidia.com, March 2022. TSMC N4, 80B transistors, 700W.
[11] NVIDIA Corporation. NVIDIA H200 SXM Product Specifications. nvidia.com, 2024. TSMC N4, 80.0B transistors (identical GH100 die as H100), 700W.
[12] NVIDIA Corporation. NVIDIA A100 Tensor Core GPU Architecture Whitepaper. nvidia.com, 2020. TSMC 7nm, 54.2B transistors, 400W.
[13] Qualcomm Technologies. Snapdragon X Elite Platform Product Brief. qualcomm.com, 2024. TSMC 4nm, 20B transistors, 80W.
[14] Smith, R. AMD FX-8150 Bulldozer Review. AnandTech, October 2011. 32nm, 1.2B transistors, 125W TDP.
[15] TechPowerUp. TechPowerUp CPU Database. techpowerup.com/cpu-specs, 2024. All historical processor specifications.
[16] Vortez. AMD Ryzen 9 9950X and 9900X Review. vortez.net, August 2024. 4nm TSMC Zen5, 20.6B total, 170W, 4.3 GHz.
[17] Vortez. AMD Ryzen 9 7950X Review. vortez.net, 2022. 5nm TSMC Zen4, 13.14B transistors, 170W TDP.
[18] Wasson, S. AMD Phenom II X4 955 Review. Tom's Hardware, April 2009. 45nm, 758M transistors, 125W TDP.

All analysis from publicly available published sources cited above. Methodology protected under Patent Applications 64/012,720 and 64/014,568. Inquiries: heath@iamperformance.net. All commercial terms through legal counsel. Heath W. Mahaffey, Independent Researcher, Entiat, WA, March 29, 2026.

