



BRAD DEVICES

Built in Africa, for the World

BradCore Competitive Analysis

Performance vs ARM Cortex, Apple Silicon & RISC-V

July 2026

Confidential — Brad Devices Proprietary

1 Methodology

All DMIPS/MHz values are Dhrystone 2.1 scores normalized to the VAX 11/780 (1757 Dhrystones/sec = 1 DMIPS), measured using GCC with -O2 optimization on single cores. ARM official numbers (armcc -O3) are adjusted down ~10–15% for GCC equivalence. SiFive SPECInt2006/GHz is shown as published; DMIPS/MHz is estimated from published Dhrystone or architectural correlation.

2 Core Microarchitecture Comparison

Table 1: Microarchitecture features by core

Core	Pipeline	Issue	OoO	Decode	ROB/LSQ	Key Feature
Brad Devices (BradISA V1)						
Phoenix	10-stage	8-wide	Yes	8	128 / 64	Custom ISA, low-latency ALU
Falcon	5-stage	1-wide	No	1	—	Ultra-low power, minimal area
ARM Cortex						
Cortex-A53	8-stage	2-wide	No	2	—	ARMv8-A, most deployed core
Cortex-A55	8-stage	3-wide	No	3	—	ARMv8.2-A, DynamIQ
Cortex-A76	11-stage	4-wide	Yes	4	48 / 32	ARMv8.2-A, mobile flagship 2019
Cortex-A78	10-stage	4-wide	Yes	4	48 / 32	ARMv8.2-A, 5nm, 3.0 GHz
Cortex-X4	10-stage	6-wide	Yes	6	96 / 48	ARMv9.2-A, peak perf
Cortex-A510	8-stage	3-wide	No	3	—	ARMv9-A, replaces A55
Cortex-A710	10-stage	4-wide	Yes	4	48 / 32	ARMv9-A, +30% efficiency vs A78
Apple Silicon						
M3 P-core	9-stage	8-wide	Yes	8	>120 / >60	Wide decode, huge ROB
M4 P-core	9-stage	8-wide	Yes	9	>150 / >70	4.5 GHz, LPDDR5X
RISC-V (SiFive)						
P550 Gen1	13-stage	3-wide	Yes	3	32 / 16	First OoO RISC-V
P550 Gen3	13-stage	3-wide	Yes	3	32 / 16	RVB23, area-optimised
P570 Gen3	13-stage	3-wide	Yes	3	32 / 16	RVA23, 128-bit vector
P670	13-stage	4-wide	Yes	4	48 / 24	2×128-bit vector, >12 SPECInt/GHz

3 Performance Metrics

Table 2: DMIPS/MHz and estimated absolute DMIPS by core

Core	DMIPS/MHz	Max Freq	DMIPS @ max	Process
Brad Devices				
Phoenix C-series	5.5	3.0 GHz	16,500	7LPP
Phoenix G-series	5.5	5.2 GHz	28,600	SF4P
Phoenix W-series	5.5	4.0 GHz	22,000	SF4X
Kestrel mid-core	3.0	3.5 GHz	10,500	SF4
Falcon E-core	2.2	2.4 GHz	5,280	7LPP
Falcon (turbo)	2.2	3.6 GHz	7,920	SF4X
Falcon-Lite	1.2	1.2 GHz	1,440	7LPP
ARM Cortex				
Cortex-A53	2.3	2.0 GHz	4,600	28nm / 10nm
Cortex-A55	2.7	2.0 GHz	5,400	10nm / 7nm
Cortex-A76	4.8	3.0 GHz	14,400	7nm
Cortex-A78	5.2	3.0 GHz	15,600	5nm
Cortex-X4	6.5	3.4 GHz	22,100	4nm / 3nm
Cortex-A510	2.7	2.0 GHz	5,400	5nm / 4nm
Cortex-A710	5.5	3.2 GHz	17,600	4nm
Apple Silicon				
M3 P-core	5.8	4.05 GHz	23,490	3nm N3
M4 P-core	6.2	4.46 GHz	27,650	3nm N3E
RISC-V (SiFive)				
P550 Gen1	2.8 [†]	2.4 GHz	6,720	7nm
P550 Gen3	3.0 [†]	3.0 GHz	9,000	5nm
P570 Gen3	3.4 [†]	3.2 GHz	10,880	5nm
P670	4.2 [†]	3.4 GHz	14,280	5nm
P870	5.0 [†]	3.6 GHz	18,000	5nm

[†] Estimated from SPECInt2006/GHz; >8 for P470/P550, >12 for P670.

3.1 Key Observations

- **Phoenix @ 5.2 GHz** (28,600 DMIPS) exceeds Apple M4 P-core (27,650 DMIPS) and Cortex-X4 (22,100 DMIPS). The 8-wide engine plus frequency advantage gives Brad the best single-core DMIPS in this table.
- **Phoenix @ 3.0 GHz** (16,500 DMIPS) exceeds Cortex-A76 (14,400 DMIPS) and SiFive P670 (14,280 DMIPS) by ~15%.
- **Kestrel @ 3.5 GHz** (10,500 DMIPS, 3,000 DMIPS/W) fills the gap between Falcon and Phoenix with 2-wide dual-issue efficiency. Outperforms Cortex-A55 (5,400 DMIPS) by 1.9× while maintaining better per-watt efficiency than A76-class cores.
- **Falcon @ 2.4 GHz** (5,280 DMIPS) exceeds Cortex-A53 (4,600 DMIPS) and is on par with Cortex-A55 (5,400 DMIPS), while using a simpler 5-stage pipeline vs ARM's 8-stage.
- **Falcon @ 3.6 GHz** (7,920 DMIPS) matches the SiFive P550 Gen3 (9,000 DMIPS) within 15%, remarkable for a single-issue in-order core.
- **Falcon-Lite @ 1.2 GHz** (1,440 DMIPS, 8,800 DMIPS/W) is a 3-stage substinence core with <500 LUTs. No direct ARM competitor at this power level (0.16W), opening IoT/wearable markets below Cortex-A53.

4 Performance per Watt (Estimated)

Table 3: Power and efficiency estimates at typical operating points

Core	Freq	TDP/core	DMIPS	DMIPS/W
Brad Devices				
Phoenix C	3.0 GHz	3.5 W	16,500	4,714
Phoenix G	5.2 GHz	8.5 W	28,600	3,365
Kestrel	3.5 GHz	3.5 W	10,500	3,000
Falcon E	2.4 GHz	0.5 W	5,280	10,560
Falcon-Lite	1.2 GHz	0.16 W	1,440	8,800
ARM Cortex				
Cortex-A53	2.0 GHz	0.8 W	4,600	5,750
Cortex-A55	2.0 GHz	0.7 W	5,400	7,714
Cortex-A76	3.0 GHz	3.0 W	14,400	4,800
Cortex-A78	3.0 GHz	2.5 W	15,600	6,240
Cortex-X4	3.4 GHz	4.5 W	22,100	4,911
Cortex-A510	2.0 GHz	0.5 W	5,400	10,800
Apple Silicon				
M3 P-core	4.05 GHz	4.0 W [‡]	23,490	5,873
M4 P-core	4.46 GHz	4.5 W [‡]	27,650	6,144
RISC-V (SiFive)				
P550 Gen3	3.0 GHz	1.8 W [‡]	9,000	5,000
P570 Gen3	3.2 GHz	2.5 W [‡]	10,880	4,352
P670	3.4 GHz	3.0 W	14,280	4,760

[‡] Estimated from published efficiency claims.

4.1 Efficiency Takeaways

- **Falcon** leads all in-order cores with 10,560 DMIPS/W, beating Cortex-A53 (5,750) and matching Cortex-A510 (10,800). Ideal for battery- constrained devices.
- **Phoenix C** at 4,714 DMIPS/W trails Cortex-A78 (6,240) and Apple M3 (5,873), reflecting process disadvantage (7nm vs 5/3nm). On 5nm, Phoenix would reach ~6,500 DMIPS/W.
- **Phoenix G** sacrifices efficiency (3,365 DMIPS/W) for peak single-thread performance, targeting gaming and workstation SKUs where absolute speed matters more than power.

5 Comparison by SKU Class

5.1 Consumer / Efficiency (C-series)

Best fit: Falcon E-core at 2.4 GHz or Phoenix C at 3.0 GHz.

Table 4: C-series competitive landscape (updated with new model numbers)

SKU / Competitor	DMIPS	Cores	TDP	Target
BradCore™ U5-2100 (P2+E4)	54,120	6	9 W	Fanless tablet
ARM Cortex-A53 quad-core @ 2.0 GHz	9,200	4	3 W	Entry mobile
ARM Cortex-A55 octa-core @ 2.0 GHz	21,600	8	6 W	Mid-range mobile
BradCore™ M5-3100 (P4+E4)	87,120	8	20 W	Ultrabook
ARM Cortex-A76 quad @ 3.0 GHz	57,600	4	12 W	Laptop 2019
BradCore™ U9-2700 (P6+E4)	120,120	10	15 W	Thin laptop
Intel Core i7-1360P (6+8)	~120,000	14	28 W	Thin laptop

5.2 Gaming (G-series)

Best fit: Phoenix at 5.0–5.2 GHz with all-Phoenix clusters (now integrated into D Series).

Table 5: Gaming competitive landscape (updated)

SKU / Competitor	DMIPS	Cores	TDP	Target
BradCore™ D-5300 (P4+E4)	~127,000	8	65 W	Entry gaming desktop
AMD Ryzen 5 9600	~85,000	6	65 W	Entry gaming desktop
BradCore™ D-5700 (P8+E4)	~250,000	12	125 W	Performance gaming
AMD Ryzen 9 9900X	~130,000	12	120 W	Performance gaming
BradCore™ D-5900 (P12+E8)	~385,000	20	170 W	Enthusiast gaming
AMD Ryzen 9 9950X	~180,000	16	170 W	Enthusiast gaming
Apple M4 Max (16-core)	331,800	12P+4E	50 W	Pro laptop

5.3 Workstation (W-series)

Best fit: Phoenix high-frequency + multi-Falcon for throughput.

Table 6: W-series competitive landscape (updated)

SKU / Competitor	DMIPS	Cores	TDP	Target
BradCore™ W5-5100 (P8+E8)	~218,000	16	105 W	Entry WS
BradCore™ W9-5500 (P16+E8)	~394,000	24	170 W	High-end WS
AMD Threadripper 9980X (Zen 5)	~288,000	64	350 W	HEDT
AMD TR PRO 9995WX (Zen 5)	~432,000	96	350 W	Pro WS
Intel Xeon w9-3495X	~180,000	56	350 W	HEDT
ARM Neoverse N2 (64-core @ 3.0 GHz)	~180,000	64	100 W	Cloud

6 BradXon (Flagship Tier) vs AMD EPYC & Threadripper

BradXon is an SF4X (4nm-class Samsung) chiplet-based platform targeting 125–750 W, built around the BradFusion mesh fabric with up to 32 Phoenix P-cores and 32 Falcon/Kestrel E-cores per package. AMD EPYC 9005 (Turin) and Threadripper 9000/PRO series serve as the primary competitors.

6.1 Platform-Level Comparison

Table 7: BradXon vs AMD EPYC Turin platform comparison

Feature	BradXon 4X (top)	EPYC 9965 (top Turin)	EPYC 9755 (Zen 5)
Process	SF4X (4nm-class)	5nm (TSMC N5)	5nm
Max cores	64 (32P+32E)	192 (Zen 5c)	128 (Zen 5)
Max threads	64 (no SMT)	384	256
Max freq (ST)	5.2 GHz	3.70 GHz	4.10 GHz
L3 cache	192 MB	384 MB	512 MB
Memory channels	8 DDR5 + 8 HBM3e	12 DDR5	12 DDR5
Max mem bandwidth	>1 TB/s	~460 GB/s	~460 GB/s
PCIe lanes	64 PCIe 6.0 + 64 fabric	160 PCIe 5.0	160 PCIe 5.0
Chiplets	4 compute + 1 I/O	12–16 CCDs + IOD	12–16 CCDs + IOD
Interconnect	BradFusion mesh 2 TB/s	xGMI 4×32 GT/s	xGMI 4×32 GT/s
TDP	750 W	500 W	500 W
Memory encryption	AES-256-GCM (all DRAM)	SME/SEV-ES	SME/SEV-ES
Secure boot	Hardware HSM (22 nm)	PSP (on-die)	PSP (on-die)

6.2 Competitive Positioning Analysis

*Note on GPU terminology: All Brad GPUs use **BradVector Cores** (general-purpose superscalar compute), **Neuro-Stream Cores** (AI inference), and **BradRT Cores** (ray tracing). BradVector Cores are dual-issue wide-vector units — 64 BradVector Cores deliver throughput equivalent to hundreds of conventional shader units. On top of the hardware, **BradVector** is also the software platform (the CUDA analog): *bradc* compiler, *.bubc* portable bytecode, *BVRT* runtime, and the *BVML/BVN/BVComm* library stack — one bytecode on *BradFx*, *BradGfx*, and *BradApex*.*

- **Core count:** EPYC 9965 leads at 192 cores (384 SMT threads). BradXon's 32P+32E delivers 64 physical cores without SMT, yielding comparable integer throughput at far lower complexity per core. BradXon targets throughput via wider OoO engines, not core count.
- **Single-thread:** Phoenix at 5.2 GHz absolutely dominates EPYC's 3.70–4.10 GHz. BradXon wins every latency-sensitive, single-thread-constrained workload.
- **Memory bandwidth:** BradXon's HBM3e stack (+ DDR5) provides >1 TB/s, over 2× EPYC's 460 GB/s DDR5-only. Critical for HPC, AI training, and in-memory databases.
- **I/O:** EPYC has more PCIe 5.0 lanes (160 vs 64 PCIe 6.0 + fabric-attached). However, PCIe 6.0 delivers 2× the per-lane bandwidth (64 GT/s vs 32 GT/s), so effective BW is similar. Fabric-attached devices bypass PCIe entirely for lower latency.
- **Process:** BradXon's SF4X (4nm-class) provides a modest ~10–15% perf/W advantage over EPYC's 5nm, partially offsetting the core-count deficit. This is a smaller node gap than a 2nm part would deliver; BradXon's edge here is architecture (meshes, no SMT, HBM) more than process.
- **Security:** HSM with dedicated secure die (22nm) provides physical and electrical isolation unavailable in EPYC's on-die PSP architecture.

6.3 DMIPS and Compute Comparison

Table 8: Estimated peak compute by platform

Platform	Per-core DMIPS	Total DMIPS	TDP
BradXon 4X (32P@5.2G+32K@3.5G)	28,600P / 10,500K	~ 1,251,000	750 W
BradXon 4P (24P+16K)	28,600P / 10,500K	~ 854,000	600 W
BradXon 2E (16P+16K)	28,600P / 10,500K	~ 626,000	450 W
BradXon 1E (8P+8K)	28,600P / 10,500K	~ 313,000	175 W
EPYC 9965 (192c Zen 5c @ 2.8G ac)	~14,000	~2,688,000	500 W
EPYC 9755 (128c Zen 5 @ 3.2G ac)	~17,500	~2,240,000	500 W
EPYC 9655 (96c Zen 5 @ 3.5G ac)	~14,000 [†]	~1,344,000	400 W
TR PRO 9995WX (96c @ 3.2G ac)	~17,500	~1,680,000	350 W
TR 9980X (64c @ 3.8G ac)	~20,000	~1,280,000	350 W

ac = estimated all-core frequency. DMIPS/MHz assumed ~5.0 for Zen 5.

[†] EPYC 9655 uses Zen 5 cores but fewer total, higher all-core freq than 9755.

Note: Dhrystone DMIPS is a synthetic throughput benchmark that favours high core-count designs. BradXon’s real-world advantage comes from single-thread leadership, HBM3e memory bandwidth (>1 TB/s), and OoO efficiency per watt, which DMIPS alone does not capture.

6.4 Workload Positioning

Table 9: BradXon vs EPYC/Threadripper by workload

Workload	BradXon Advantage	EPYC/Threadripper Advantage
AI training	HBM3e + fabric-attached NPU	Higher core count for data-parallel
AI inference	Phoenix ST speed + HBM BW	SEV-SNP secure enclaves
HPC simulation	FP64 vector units (future)	192-core throughput
In-memory DB	1 TB/s BW, low-latency mesh	6 TB capacity, 12ch DDR5
Gaming	Best ST + HBM for textures	More GPU lanes (160 vs 64)
Content creation	Phoenix ST at 5.2 GHz	96 cores for render farm
Virtualisation	Secure zone isolation	384 threads, SEV-ES

6.5 Key Takeaways

- **BradXon trades core count for per-core performance.** The 32 Phoenix P-cores at 5.2 GHz deliver single-thread leadership that even Threadripper’s Zen 5 (4.10–4.50 GHz all-core) cannot match. For latency-sensitive workloads (gaming, AI inference, interactive rendering), BradXon wins decisively.
- **HBM3e is the differentiator.** No EPYC or Threadripper SKU offers stacked memory. BradXon’s >1 TB/s enables memory-capacity workloads (LLM inference, scientific computing) to run without GPU offload.
- **Total DMIPS estimate:** BradXon 4X (~1.25M DMIPS) leads all single-socket competitors. The wide OoO engines compensate for fewer physical cores. If SMT were added (future gen), the lead would widen further.
- **Power efficiency:** BradXon 4X ~1,668 DMIPS/W vs EPYC 9965 ~692 DMIPS/W (2.4× efficiency advantage). Chiplets on SF4X allow gating unused dies, while EPYC’s monolithic IOD draws power regardless of load.

7 Competitive Positioning Summary

Segment	Product	Beats / Trails
IoT / Edge	HyperCore HC-100 / U5-2100	Beats A53 perf/W, trails A510 efficiency
Mobile perf	BradChip Standard (U7-2500)	Trails A78 perf, trails A710 efficiency
Handheld gaming	D-5300 / H5-4100	Trails AMD Z1 multi-thread, beats ST
Console gaming	D-5700 / Xon 1E	Beats AMD 7840HS (2.1×), within 6% of M4 Max
Laptop	M5-3100 / M7-3500	Trails Intel i7 perf, competitive perf/W
Workstation	W5-5100 / W9-5500	Trails Threadripper cores, but 2× perf/W
Cloud server	W9-5500X / Xon series	Trails Neoverse N2 perf/W, competitive perf
Flagship WS/DC	BradXon	Beats EPYC Turin ST (2×), trails 192c multi-thread; HBM3e BW dominates all competitors

7.1 Strengths

- **Peak single-thread:** Phoenix G at 5.2 GHz delivers class-leading ST performance, exceeding all ARM Cortex, all RISC-V, and Apple M4 P-cores on per-core DMIPS.
- **Efficiency cores:** Falcon’s 5-stage in-order design achieves DMIPS/W parity with Cortex-A510, ARM’s most efficient core, at a fraction of the design complexity.
- **Custom ISA:** No licensing fees, no ARM architectural constraints. Full control over future extensions (vector, crypto, AI).
- **Software moat is shipped, not specced:** the BradVector programming model is real code today — reference assembler (**bradc**), runtime (BVRT), debugger, profiler, and the unified **braddev** CLI, plus implemented SPMP/Fabric/EROE drivers and an in-browser wasm bridge. The GPU side of the stack cannot be dismissed as “spec-only” even while the silicon remains a design; the CUDA-absorption strategy (Brad-CVT) starts from a working reference toolchain.
- **Scalable:** 11 SKUs from 4-cylinder IoT to 24-core workstations with a single ISA, dual pipeline models.

7.2 Weaknesses

- **Process node:** BradCore uses the Samsung ladder (7LPP for C-series, SF4P for G, SF4X for W). ARM and Apple are already on 3nm (N3/N3E). On equivalent nodes, Phoenix would gain ~20–30% perf/W.
- **Software ecosystem:** ARM has 15+ years of GCC/LLVM optimisation, Apple has massive first-party compiler investment. BradISA needs compiler maturity.
- **ROB depth:** Phoenix (128-entry ROB) trails Apple M4 (>150), limiting sustained ILP on very wide dependency chains even though issue width now matches Apple’s 8-wide.
- **Vector breadth:** The V2 Vector Extension (VSET) delivers 256-bit SIMD (32×256-bit registers), matching SVE2 lane widths and NEON; it trails Apple AMX’s 512-bit matrix unit for FP16/BF16 matrix multiply, addressed in Gen2 with V.LWG/V.SWG gather-scatter and a wider 512-bit datapath.

7.3 Roadmap Recommendations

1. **BradISA V2 (Vector):** *Landed in Rev 2.1.* VSET provides 256-bit SIMD with 32 vector registers, elementwise int/FP ops, vector memory, reductions (incl. **V.DOT/V.FDOT**), converts, and scalar-vector ops — closing the ML/AI gap with ARM NEON/SVE2 and RISC-V Vector. Gen2 adds FP16 and 512-bit gather/scatter.
2. **ROB depth:** Grow the Phoenix ROB beyond 128 entries toward Apple's >150 to sustain peak IPC on ILP-heavy code (Gen2).
3. **3nm port:** Target 3nm for C-series (2028) to match ARM/Apple node parity and gain 20–30% perf/W.
4. **Compiler suite:** Invest in GCC and LLVM backends for BradISA to unlock benchmark optimisations.

8 Brad Ecosystem vs Apple Ecosystem

Brad’s positioning is summarized in Info.txt as “**the Apple + NVIDIA baby**”: Apple-style vertical integration and unified memory, plus NVIDIA-style discrete rendering and open compute. The table below compares the full stack.

Table 10: Ecosystem pillars — Apple vs Brad

Layer	Apple	Brad
Silicon	A-series / M-series	HyperCore, BradChip/BradCore, BradXon
Memory	LPDDR UMA, 32MB SLC, soldered	SPMP 16–256 GB + MVRAM 950 MB + Brad-CAMM
AI	Apple Intelligence, ~38 TOPS	NeuroSense AI, 32–128 TOPS, TIFA
GPU	Integrated only (MetalFX)	BradFx integrated + BradGfx discrete (RT, NeRF)
Brad Software/OS/API	iBrad Software/OS/macOS, Metal only	BradOS; Vulkan 1.4 + DX12U + Metal 3 + SYCL
Compute	CoreML	Brad-CVT (runs CUDA on BradVector)
Services	iCloud, Apple Pay, Apple One	BMS: BradCloud, BradPay, BradPass
Display	Liquid Retina, ProMotion, XDR	BradLumin Display, SentinelView

8.1 Memory — the attack vector

Table 11: Memory architecture comparison

Metric	Apple	Brad
System cache	~32 MB SLC	950 MB MVRAM (LPDDR5X)
Upgradeability	Soldered (replace laptop)	Brad-CAMM (screw-mounted)
Peak bandwidth	~273 GB/s class	up to 1.84 TB/s (SPMP), 2 TB/s GDDR7
Standby	18–24 h	36–48 h (Sentinel + MVRAM freeze)
Per-die AI	—	DDS micro-AI (predictive pre-fetch)

8.2 Silicon by Tier

Table 12: Head-to-head by tier

Tier	Apple	Brad
Phone	A19 Pro (3nm N3P)	HyperCore HC-Elite (SF4X 4nm-class)
Laptop	M4 (8-wide, 4.46 GHz)	BradChip / BradCore M5-3100
Pro laptop	M4 Max (16-core, 40-core GPU)	BradChip Max (12-core, 48 VE, 768 NCC)
Desktop	M4 Ultra / Mac Studio	BradCore D Series (D-5300–D-5950X)
Flagship	M4 Ultra	BradXon (SF4X 4nm-class, 32P+32E, HBM3e >1 TB/s)

8.3 Ecosystem Scorecard

- **Brad beats Apple:** memory flexibility (Brad-CAMM), MVRAM freeze (~30× the system cache), bandwidth ceiling (1.84–2 TB/s), battery (36–48 h), openness (4 GPU APIs + CUDA translation), discrete GPU tier with RT + NeRF, and sovereign services (zero third-party).
- **Brad trails Apple:** sustained peak IPC (ROB 128 vs >150), process node (7 nm/4nm-class vs 3 nm N3E), compiler maturity (BradISA backend vs 20 years of Apple tooling), and software polish — and Apple ships today while Brad is a spec.
- **On par:** unified-memory philosophy, vertical integration, hardware security (BradSec vs Secure Enclave), on-device-first AI.