



BRAD DEVICES

Built in Africa, for the World

Brad Devices Product Portfolio

Complete Product Lineup — Desktop, Mobile, Graphics

Revision 1.0 — July 2026

Gen 1 Kinetic (2026) | Gen 2 Synaptic (2027–28) | Gen 3 Quantum (2029+)

Confidential — Brad Devices Proprietary

Contents

1	Product Line Architecture	3
1.1	GPU Core Architecture	3
1.2	Generation Naming	4
2	Desktop Processors — BradCore D Series	5
2.1	D Series SKU Lineup	5
2.2	Key Features	5
2.3	Competitive Positioning	5
3	Mobile Processors — BradCore U, M, H Series	6
3.1	U Series — Ultra-Low Power (9–15W)	6
3.2	M Series — Mainstream Mobile (15–35W)	6
3.3	H Series — High-Performance Mobile (35–65W)	6
4	Workstation Processors — BradCore W Series	7
5	BradXon — Flagship Chiplet Series	8
6	BradGfx — Discrete Graphics	9
6.1	Core Architecture	9
6.2	Gen1 BradGfx Lineup	9
6.3	Specifications	9
6.4	Competitive Positioning	10
6.5	Key Technologies	10
7	BradFx — Integrated Graphics	11
7.1	BradFx Lineup	11
7.2	Specialised Hardware	11
7.3	Performance Estimates	11
8	BradChip — Unified SoC Packages	12
9	HyperCore — Mobile NPU SoC	13
10	Target Device-to-Product Mapping	14
11	Product Model Numbering Convention	15
11.1	CPU Model Numbers	15
11.2	GPU Model Numbers	15
12	Software & Toolchain Portfolio (Shipped Reference Implementations)	15
13	Revision History	16

1 Product Line Architecture




Figure 1: Product brand hierarchy — full catalog in the table below

Brand	Product Line	Description
BradCore	D Series	Desktop flagship (socketed LGA, P-core heavy)
	U Series	Ultra-low-power mobile (9–15W BGA)
	M Series	Mainstream mobile (15–35W BGA)
	H Series	High-performance mobile (35–65W BGA)
	W Series	Workstation (65–150W LGA/BGA)
	Xon 1–4P/E/X	Flagship chiplet CPU (125–750W, SF4X 4nm-class)
BradChip	Lite–Max	Unified SoC (CPU+GPU+NPU, Matrix socket)
	100 / 300 / 500 / 700	Discrete GPU (75–350W)
	Max (i)	Flagship discrete/integrated (600W)
BradGfx	10 / 30 / 50 / 70 / 90	Integrated GPU (5–55W)
HyperCore	HC-100–Elite	Mobile SoC (NPU-fused, 7LPP/SF4/SF4P)
BradFx		

Figure 2: Brad Devices complete product architecture

1.1 GPU Core Architecture

All Brad GPUs — discrete (BradGfx) and integrated (BradFx) — are built from three specialised processor arrays:

- **BradVector Cores:** General-purpose compute engines. Dual-issue, superscalar, wide-vector design. Each core handles vector FMA + scalar ALU + branch per cycle, delivering the throughput of 16–32 conventional shader units. Rasterization, vertex/pixel shaders, physics, and general compute all run on BradVector Cores.
- **Neuro-Stream Cores (BradGfx) / Neuro-Lite Cores (BradFx):** On-GPU AI accelerators for BradSense-Render (upscaling, frame generation), neural denoising, and real-time inference.
- **BradRT Cores:** Dedicated ray tracing units with BVH traversal, ray-primitive intersection, NeRF evaluation, and programmable ray pipelining.

1.2 Generation Naming

Generation	Years	Process Nodes
Kinetic (Gen 1)	2026	7LPP (budget), SF4/SF4P (mainstream), SF4X (flagship)
Synaptic (Gen 2)	2027–2028	SF3/SF2, 1.4nm
Quantum (Gen 3)	2029+	Sub-1nm, graphene channel

2 Desktop Processors — BradCore D Series

The D series is BradCore’s desktop flagship line, designed for socketed LGA motherboards (BradBoard Nexus). D chips use P-core-dominant configurations for maximum single-thread and multi-thread throughput.

2.1 D Series SKU Lineup

Table 1: BradCore D Series — Desktop Processors

Model	P-cores	E-cores	L3 MB	Base/Boost GHz	TDP W	PCIe
D-5300	4	4	16	3.4 / 4.8	65	20 lanes PCIe 6.0
D-5400	4	8	24	3.4 / 4.8	65	20 lanes PCIe 6.0
D-5500	6	4	32	3.6 / 5.0	95	24 lanes PCIe 6.0
D-5600	6	8	40	3.6 / 5.0	95	24 lanes PCIe 6.0
D-5700	8	4	48	3.8 / 5.2	125	28 lanes PCIe 6.0
D-5800	8	8	56	3.8 / 5.2	125	28 lanes PCIe 6.0
D-5900	12	8	80	3.6 / 5.2	170	28 lanes PCIe 6.0
D-5950X	16	8	96	3.4 / 5.0	230	28 lanes PCIe 6.0

2.2 Key Features

- All D SKUs use Phoenix P-cores (5.5 DMIPS/MHz) with BDX optimizations
- Socket: LGA3 (BradBoard Nexus Pro, + Matrix Studio via adapter)
- Memory: 2–4 channel Brad-CAMM DDR5/DDR6 (up to 70 GB/s per channel)
- Fabric: BradFusion mesh (on-die) + BradLink fabric bridge for multi-socket
- BradSec: BradSec-RAM on all SKUs, BradSec-X on D-5900+
- NPU: Integrated NeuroSense AI Engine (384–1024 NCC depending on SKU)
- Process: SF4P (D-5300–5600), SF4X (D-5700+)

2.3 Competitive Positioning

Table 2: D Series vs Competitors

SKU	Competitor	Total DMIPS	Target
D-5300	Ryzen 5 9600 / Core i5-14600	~127,000	Budget desktop
D-5500	Ryzen 7 9700X / Core i5-15600	~186,000	Mainstream desktop
D-5700	Ryzen 9 9900X / Core i7-15700	~250,000	Performance desktop
D-5900	Ryzen 9 9950X / Core i9-15900K	~385,000	Enthusiast desktop
D-5950X	Threadripper 7960X / Core Ultra 9 285K	~482,000	Extreme desktop

3 Mobile Processors — BradCore U, M, H Series

3.1 U Series — Ultra-Low Power (9–15W)

For fanless thin-and-light laptops, tablets, and foldables.

Table 3: BradCore U Series — Ultra-Low Power Mobile

Model	P-cores	E-cores	L3 MB	Freq GHz	TDP W	GPU
U5-2100	2	4	8	2.0 / 3.8	9	BradFx 10
U5-2300	4	4	12	2.0 / 4.0	12	BradFx 10
U7-2500	4	4	16	2.4 / 4.2	15	BradFx 30
U9-2700	6	4	24	2.6 / 4.4	15	BradFx 30

Target: BradMatrix Essential, BradSlate, BradPhone Ultra (HyperCore on phone SoC).

3.2 M Series — Mainstream Mobile (15–35W)

For thin laptops, ultrabooks, and productivity devices.

Table 4: BradCore M Series — Mainstream Mobile

Model	P-cores	E-cores	L3 MB	Base/Boost GHz	TDP W	GPU
M5-3100	4	4	16	2.4 / 4.4	20	BradFx 30
M5-3300	4	8	24	2.4 / 4.4	25	BradFx 50
M7-3500	6	4	32	2.8 / 4.8	28	BradFx 50
M7-3700	6	8	40	2.8 / 4.8	35	BradFx 70
M9-3900	8	4	48	3.0 / 5.0	35	BradFx 70

Target: BradMatrix Prime 14", BradSlate Ultra.

3.3 H Series — High-Performance Mobile (35–65W)

For gaming laptops, mobile workstations, and creator devices.

Table 5: BradCore H Series — High-Performance Mobile

Model	P-cores	E-cores	L3 MB	Base/Boost GHz	TDP W	GPU
H5-4100	6	4	40	3.0 / 5.0	45	BradFx 90
H7-4300	8	4	56	3.2 / 5.2	55	BradFx 90
H9-4500	8	8	64	3.2 / 5.2	65	BradFx 90 + BradGfx 300M
H9-4700X	12	8	80	3.0 / 5.0	65	BradFx 90 + BradGfx 500M

Target: BradMatrix Ultra 16", BradOS Compute Mobile.

4 Workstation Processors — BradCore W Series

For professional workstations, render farms, and server-class single-socket systems.

Table 6: BradCore W Series — Workstation

Model	P-cores	E-cores	L3 MB	Base/Boost GHz	TDP W	Memory
W5-5100	8	8	64	3.0 / 4.6	105	4ch DDR5 / DDR6
W7-5300	12	8	96	3.2 / 4.8	125	4ch DDR5 / DDR6
W9-5500	16	8	128	3.4 / 5.0	170	8ch DDR5 / DDR6
W9-5700X	16	16	192	3.2 / 4.8	200	8ch DDR6

Features:

- All W-series support ECC memory
- BradSec-X (W9+) or BradSec-RAM (W5/W7)
- BradBoard Nexus Pro or Matrix Studio
- Optional BradGfx 300–700 discrete GPU

5 BradXon — Flagship Chiplet Series

Table 7: BradXon lineup (SF4X 4nm-class, BradFusion mesh)

SKU	Chiplets	P-cores	E-cores	L3 MB	TDP W	Config
Xon 1P	1	4	0	32	125	Single, compact WS
Xon 1E	1	8	8	48	175	Single, balanced
Xon 2P	2	12	0	80	350	Dual, P-only
Xon 2E	2	16	16	96	450	Dual, balanced
Xon 4P	4	24	0	160	600	Quad, P-only
Xon 4X	4	32	32	192	750	Quad, extreme

Memory: HBM3e (2–8 stacks, 128–512 GB/s each) + optional DDR5. **Fabric:** BradFusion mesh, 2 TB/s bisection bandwidth (4-die). **Target:** BradOS Compute Pro, BradAstra servers, BradSkynet (controller). BradAxis uses standard BradCore + GDDR7 (see Console spec).

6 BradGfx — Discrete Graphics

BradGfx is the discrete GPU family, available as standalone PCIe 6.0 cards or embedded via ANF (Adaptive Neural Fabric) on BradXon. Every BradGfx GPU is built from three specialised processor arrays: **BradVector Cores** (general-purpose compute), **Neuro-Stream Cores** (AI inference), and **BradRT Cores** (ray tracing).

6.1 Core Architecture

BradVector Cores are dual-issue, superscalar, wide-vector compute engines. Unlike traditional SIMD shaders that execute one narrow operation per cycle, each BradVector Core processes multiple operations—a vector FMA, a scalar ALU op, and a branch—in a single clock cycle. A small number of high-IPC BradVector Cores delivers the throughput of thousands of conventional shader cores, with dramatically better power efficiency and die utilisation.

Neuro-Stream Cores are the on-GPU AI accelerators, handling BradSense-Render (AI upscaling, frame generation), neural denoising, MoE expert switching, and real-time physics inference. They complement the system BradNPU, which handles heavier sustained AI workloads. The cores are GEN 1 design; the BradSense-Render algorithm behind them is not waiting on them — a software reference upscaler already ships on the BradVector reference stack (spatial core, PSNR-baselined and unit-tested), with the temporal core queued.

BradRT Cores provide dedicated hardware for BVH traversal, ray-primitive intersection, NeRF evaluation, and programmable ray pipelining—enabling real-time path tracing and global illumination.

6.2 Gen1 BradGfx Lineup

Table 8: BradGfx Discrete GPU SKUs

Model	Tier	BradVector	BradRT	Neuro-Stream	VRAM	TDP
BradGfx 100	Lite	64	16	16	8 GB GDDR7	75W
BradGfx 300	Standard	128	32	32	12 GB GDDR7	130W
BradGfx 500	Ultra	256	64	64	16 GB GDDR7	220W
BradGfx 700	Max	512	128	128	24 GB GDDR7	350W
BradGfx Max	ANF	1024	256	192	48 GB HBM3e	600W
BradGfx Max Xtensive	Xtensive	256	128	512	32 GB GDDR7	225W

6.3 Specifications

Table 9: BradGfx detailed specifications

Model	FP32 TFLOPS	RT GigaRays/s	L2 MB	Mem BW	Form
100	4–6	4	2	256 GB/s	PCIe 6.0 x8
300	10–14	10	4	512 GB/s	PCIe 6.0 x16
500	20–28	20	6	768 GB/s	PCIe 6.0 x16
700	40–55	40+	8	1.0 TB/s	PCIe 6.0 x16
Max	80–110	80+	8 + 2 GB IB	2.0 TB/s	Fabric / PCIe 6.0 x16
Max Xtensive	80–110	40+	8 + 2 GB IB	0.75–2.0 TB/s + SPMP	PCIe 6.0 x16

6.4 Competitive Positioning

Table 10: BradGfx vs NVIDIA and AMD

BradGfx	NVIDIA Class	AMD Class
100	RTX 4060	RX 7600
300	RTX 4070	RX 7800 XT
500	RTX 4080	RX 7900 XT
700	RTX 5080	RX 9070 XT
Max	RTX 5090 Ti / Titan	MI350X
Max Xtensive	RTX 5090	MI350X

6.5 Key Technologies

- **BradVector Cores:** Dual-issue superscalar, wide-vector; each core handles vector FMA + scalar ALU + branch per cycle. A single BradVector Core replaces 16–32 conventional shader units.
- **BradRT Cores:** 2nd-gen ray tracing with BVH traversal, triangle-cluster accelerator, NeRF evaluation, programmable ray pipe.
- **Neuro-Stream Cores:** On-GPU AI for BradSense-Render (upscaling + frame gen), AI denoising, MoE expert switching.
- **Infinity Buffer:** On-die SRAM LLC (BradGfx 700+, Max); reduces SPMP traffic by ~40% for frame-cache workloads.
- **Hybrid-UMA (Max Xtensive):** Dedicated GDDR7 VRAM on a 192–512-bit bus (0.75–2.0 TB/s) for the Windows/eGPU role, with zero-copy SPMP overflow on BradOS — one pool, two tiers.
- **VRS:** Variable-rate shading tier 2 hardware.
- **BradSec Display:** Pixel-level encryption, HDCP 2.3, secure framebuffer path.

7 BradFx — Integrated Graphics

BradFx is the integrated GPU family, sharing die area with BradCore on SoC packages. All BradFx SKUs use the exact same **BradVector Core** architecture found in BradGfx, scaled for power efficiency. They also include dedicated **Neuro-Lite Cores** (subsets of the full Neuro-Stream design) and **BradRT-Lite** ray intersection units, ensuring consistent ISA across the entire Brad GPU ecosystem.

7.1 BradFx Lineup

Table 11: BradFx Integrated GPU SKUs

Model	Tier	BradVector	BradRT	Neuro-Lite	Freq	TDP	In
Fx 10	Lite	8	—	—	800 MHz	5W	U5 series, HyperCore Lite
Fx 30	Standard	16	—	2	1.0 GHz	12W	U7+, M5, BradChip Std
Fx 50	Ultra	32	8	4	1.2 GHz	20W	M7, BradChip Ultra
Fx 70	Max	48	16	8	1.4 GHz	35W	M9, BradChip Elite
Fx 90	Max-i	64	24	12	1.6 GHz	55W	H5+, BradChip Max

7.2 Specialised Hardware

- **BradVision Engine:** Dedicated 8K AV1 / BradCompress video decode/encode.
- **Neuro-Lite Cores:** On-GPU AI for real-time effects (background blur, low-res upscaling, frame synthesis) without touching the main NPU.
- **BradRT-Lite:** Ray intersection hardware on Fx 50+, enabling UI glass effects and real-time shadows.
- **Infinity-Tile SRAM:** 16–32 MB on-die buffer (Fx 50+) reducing memory system pressure by ~40%.

7.3 Performance Estimates

Table 12: BradFx FP32 TFLOPS and competitive class

Model	FP32 TFLOPS	Competitive Class
Fx 10	1–2	Intel UHD Graphics
Fx 30	2–4	Apple M2 GPU / RDNA 2 mobile
Fx 50	5–8	Apple M3 Pro GPU
Fx 70	10–15	Apple M4 Max GPU / RTX 3050 mobile
Fx 90	15–22	RTX 4060 mobile

8 BradChip — Unified SoC Packages

BradChip combines BradCore CPU + BradFx GPU + BradNPU into a single package that sockets into BradBoard Matrix motherboards.

Table 13: BradChip tier specifications

Tier	CPU	GPU	NPU	Process	Target
Lite	P2+E4	Fx 10	256 NCC	7LPP	BradMatrix Essential
Standard	P4+E4	Fx 30	384 NCC	SF4	BradMatrix Base, BradSlate Prime
Ultra	P6+E4	Fx 50	512 NCC	SF4P	BradMatrix Prime 14"
Elite	P8+E4	Fx 70	768 NCC	SF4X	BradMatrix Ultra 16"
Max	P8+E8	Fx 90	1024 NCC	SF4X	BradOS Compute Base

9 HyperCore — Mobile NPU SoC

HyperCore is Brad’s mobile-first SoC for BradPhone and BradSlate, integrating CPU + GPU + massive NPU arrays for on-device AI.

Table 14: HyperCore Atlas SKU Lineup

SKU	CPU	GPU	NPU Cores	MVRAM
HC-100 (Essential)	P1+E4+S	BradFx 10	16	128 MB
HC-300 (Base)	P2+E4+S	BradFx 30	32	256 MB
HC-500 (Ultra)	P4+E4+S	BradFx 50	48	450 MB
HC-Elite (Prime)	P6+E4+S	BradFx 70	64	950 MB

Key: S = Sentinel always-on core (microwatt-class). MVRAM = LPDDR5X on-die buffer. LPGDDR = optional low-rate GDDR6/7 graphics tier, strictly mobile (HyperCore/Atlas BradPhone GPUs); desktop BradGfx uses standard JEDEC GDDR7. Process: Samsung tiered ladder (7LPP/SF4/SF4P/SF4X).

10 Target Device-to-Product Mapping

Table 15: Complete product-to-device mapping

Device	SoC/CPU	GPU	Board
BradPhone Essential	HyperCore HC-100	BradFx 10	Mobile Nano
BradPhone Base	HyperCore HC-300	BradFx 30	Mobile Standard
BradPhone Ultra	HyperCore HC-500	BradFx 50	Mobile Elite
BradSlate Base	HyperCore HC-300	BradFx 30	Mobile Standard
BradSlate Prime	HyperCore HC-Elite	BradFx 70	Mobile Elite
BradSlate Ultra	BradCore M9-3900	BradFx 70	Mobile Elite
BradMatrix Essential	BradChip Lite	BradFx 10	Matrix Air
BradMatrix Base	BradChip Standard	BradFx 30	Matrix Air
BradMatrix Prime 14"	BradChip Ultra	BradFx 50	Matrix Pro
BradMatrix Ultra 16"	BradChip Elite	BradFx 70	Matrix Pro
BradOS Compute	BradChip Max	BradFx 90	Matrix Studio
BradOS Compute Pro	BradCore D-5900+	BradGfx 500+	Nexus Pro
BradCube Base	BradChip Max	BradFx 90 (Max-i)	Matrix Studio-class
BradCube Prime	BradChip Max (bin)	BradFx 90 + eGPU	Matrix Studio-class
BradCube X Prime	BradXon (single)	BradGfx	Studio X-class
BradCube X Ultra	BradXon (1-2)	BradGfx Max	Studio X Ultra-class
BradAxis Console	BradCore console SKU (8P+8E)	BradGfx Max (fabric)	Nexus Console
BradAstra Headset	BradNeuro X (XR SKU)	BradFx 90 (integrated)	Mobile Elite
BradAstra Server	BradXon Xon 4X	BradGfx Max ×4	Rack Quad
BradSkynet Node	BradApex + Xon 2P (ctrl)	— (AIMP)	Rack Sled

11 Product Model Numbering Convention

11.1 CPU Model Numbers

BradCore {Family}{Gen}{SKU}{Suffix}

- **Family:** D (desktop), U/M/H (mobile), W (workstation)
- **Gen:** 1st digit = generation (2 = Gen1 Kinetic, 3 = Gen2 Synaptic, 4 = Gen3 Quantum)
- **SKU:** Next 3 digits = performance tier (100–900)
- **Suffix:** X = extreme, X3D = 3D V-cache, U = ultra-low power, HX = unlocked

Examples:

- D-5700 = Desktop, Gen1, 700-tier performance (8 cores)
- M7-3500 = Mobile M-series, Gen1, 500-tier (6 cores)
- U5-2100 = Mobile U-series, Gen1, 100-tier (2+4 cores)
- W9-5500X = Workstation, Gen1, 500-tier extreme (16+8 cores)

11.2 GPU Model Numbers

BradGfx {Hundreds} {Suffix}

BradFx {Tens} {Suffix}

- BradGfx: 100 (entry) → 300 → 500 → 700 → Max (integrated flagship) → Max Xtensive (GDDR7 hybrid-UMA edition)
- BradFx: 10 (entry) → 30 → 50 → 70 → 90 (flagship)
- Suffix: M = mobile, Ti = enhanced, X = extreme clock
- Max Xtensive: flagship Max die on GDDR7 192–512-bit bus, hybrid-UMA (dedicated VRAM + SPMP overflow)

12 Software & Toolchain Portfolio (Shipped Reference Implementations)

The SKU line is silicon-centric; the reference *software* stack that makes the line programmable is real code shipping today in the repository — no paper APIs:

- **Reference drivers** — SPMP (first-fit pool allocator, coalescing, fragmentation stats), Fabric (DMA / atomics / bind-memory), EROE (thermal + fan-curve + workload governor) in `src/drivers/`, exercised by the driver test suite.
- **bradlib Platform API** — one C surface for the whole BradVector platform: compile → pack → session → launch → run → inspect → breakpoints (`src/include/brad/bradlib.h`).
- **In-browser toolchain** — the 26-export wasm bridge mirrored 1:1 by `site/js/bradvector.js`; the site assembles and runs kernels live.
- **BMS services** — REST v1 spec + stdlib-only reference server (`src/bms/`), all request/response cases verified.
- **TIFA-SDK** — Python-first AI toolchain (`import tifa`) with exact FP8/FP4/INT4 digit formats (bit-identical spot checks).
- **BradDev Kit** — the one-binary developer CLI `braddev` (`cvt`, `run`, `dbg`, `gpu`, `timeline`, `test`); SAXPY and fib kernels verified end-to-end with zero traps.

13 Revision History

Revision	Date	Changes
1.0	July 2026	Initial portfolio definition: D/U/M/H desktop & mobile series, BradGfx 100–700 / BradFx 10–90 numbered SKUs, BradChip Lite–Max tiers, HyperCore lineup
1.1	Aug 2026	Added BradGfx Max Xtensive: 256 BradVector / 128 BradRT / 512 Neuro-Stream, 32 GB GDDR7 on 192–512-bit bus, hybrid-UMA (dedicated VRAM + SPMP overflow), 225W
1.2	Sept 2026	Software receipts section: reference drivers, bradlib Platform API, wasm bridge, BMS server, TIFA-SDK, and the unified BradDev Kit CLI — all shipped code