

Synthetic Brain Variant D: Two Builds, 2026 and 2030

Oct 8, 2026 · @David N Sutton

Variant D is Variant C specified twice: a build that can be ordered today, and a build that depends on research now in progress. Same architecture, same module boundaries, two sets of parts. The point is that the 2026 build is not a prototype of the 2030 one — it is the same brain with slower organs.

The D rule

Variant C named one technology per brain area. Variant D names two, and holds the interfaces between modules fixed so either can be dropped in.

Column	Definition	Test for inclusion
2026 build	Parts purchasable now, or published with specifications	Could be quoted by a vendor this quarter
2030 build	Research lines with a public target inside the window	A named group has stated an aim, with a working result behind it

What is deliberately excluded from the 2030 column. Anything with no demonstrated precursor. A photonic processor that is small today but real goes in; a room-temperature fault-tolerant quantum computer with a million qubits does not, because nothing in the public record supports that date. Where a bet is weak, the section on the 2030 bets says so rather than quietly dropping it.

The interface contract. Three things stay identical between builds, and this is what makes D a single design rather than two:

- 1. Module boundaries.** The same 14 modules, the same job each.
- 2. The four buses.** Motor fieldbus, optical backbone, hard-wired safety line, and the four-scalar state bus. A 2030 module must speak the same protocol as the 2026 one it replaces.
- 3. The override rule.** Faster, lower layers override slower, higher ones. No 2030 part is allowed to make a reflex depend on a planner.

Because of that contract, the build is not two projects. It is one robot in which modules are swapped as their 2030 parts arrive, and every swap is independently testable against the module it replaces.

Area by area

Seventeen modules, each specified twice. Rows are ordered from the fastest layer up to the slowest, as in the base design.

Module	2026 build	2030 build	What the upgrade buys
Spinal cord	Microcontroller plus FPGA at each joint, 10–20 kHz current loop	Same silicon, integrated into the actuator as one smart-joint package	Less wiring and mass, not more capability. This module is already solved
Brainstem	Lockstep safety microcontroller, hard-wired emergency stop	Same, with a formally verified core	Certification evidence rather than new function
Cerebellum	FPGA adaptive filters, one per joint, 1 kHz	Analog in-memory crossbar array running the forward model at the same rate	10–100x less power for the same correction; whole-body model instead of per-joint
Basal ganglia	Reinforcement-learning policy on CPU plus hardware winner-take-all gate	Same policy on a photonic inference core, with the arbitration gate still electronic	Faster action selection; the gate stays electronic on purpose
Thalamus	Time-sensitive Ethernet switch with priority gating	Optical circuit switch with wavelength-addressed gating, no conversion	Removes every optical-to-electronic hop in the sensory path

Module	2026 build	2030 build	What the upgrade buys
Sensory cortex	NPU and GPU perception models, 30–60 Hz vision	Integrated photonic neural network taking camera light directly	Perception at light speed, and no analog-to-digital stage at the front
Association cortex	World model on GPU, 10–30 Hz	Photonic matrix cores with phase-change synapses for in-place learning	The world model updates while running instead of between sessions
Prefrontal cortex	Large vision-language model on GPU, 1–10 Hz	Hybrid photonic-electronic inference, electronic memory retained	Faster planning; memory bandwidth is still the binding constraint
Motor cortex	Learned motion policy on GPU, 100 Hz	Photonic policy inference feeding the same fieldbus	Lower latency into the cerebellum loop
Hippocampus	Hippocampal assembloid plus electrode array, electronic MIMO twin, diamond archive	Perfused, vascularised assembloid with full optical read and write, running for years	The memory unit stops being a consumable. This is the single biggest change in D
Entorhinal cortex	Grid-cell model in software, theta paced electrically	Grown as part of the assembloid, generating its own theta	Native rhythm instead of an imposed one
Amygdala	Fast classifier with a hard line to the brainstem	Same, on a dedicated low-power inference chip	Always-on at negligible power
Hypothalamus	Embedded controller for battery, thermal, and culture life support	Same, plus closed-loop chemistry control of the wet modules	Life support stops needing a human in the loop

Module	2026 build	2030 build	What the upgrade buys
Neuromodulator nuclei	Four scalars on a broadcast bus; microfluidic dopamine to the tissue	Same bus; photo-uncaging delivers transmitters with millisecond timing	Chemical signals get the timing precision electrical ones already have
SCN	Disciplined oscillator plus PTP time master	Chip-scale optical clock	Better long-term stability; marginal for a robot
White matter	Fiber with wavelength multiplexing, plus copper for safety lines	Same fiber, denser multiplexing, co-packaged optics at every module	Bandwidth stops being a design constraint anywhere
Quantum co-processor	D-Wave annealer over the cloud, for planner optimisation only	Early fault-tolerant machine, still planner-only	Sampling for planning under uncertainty; no brain area moves onto it


Three rows where the 2026 and 2030 builds are nearly the same, and that is the correct answer rather than a gap: spinal cord, brainstem and amygdala. These are the safety-critical, latency-critical modules. Silicon already does their job, and the 2030 column adds verification and packaging, not capability. A design that proposed replacing them with something newer would be worse, not better.

Which upgrades to plan around

A dated timeline would be fiction. Only two of the ten bets have a public target date at all, and the rest would need invented ones. What can be said honestly is which upgrades deserve to shape the design now.

Three upgrades are worth planning around; two need a fallback

Vertical: impact on the system if it lands. Horizontal: confidence it arrives by 2030.



2030 upgrades - 10 bets on two axes

Plan on these three. They either already exist in another industry or have a working result behind them. Build the 2026 version so the swap is a drop-in.

Hedge these two. The vascularised assembloid and scaled phase-change synapses are the high-impact, low-confidence pair, and both need their fallback engineered into the 2026 build rather than bolted on later. For the assembloid that means the staggered-culture scheme and the electronic twin; for the synapses it means keeping the GPU learning path alive.

The bottom-left group is not a criticism. A fault-tolerant quantum co-processor may well arrive. It sits there because in this architecture it only ever accelerates the planner, so its arrival changes little about the brain.

The 2030 bets

Every entry in the 2030 column is a bet on a research line. Some are near-certain, some are not, and the design should not pretend otherwise. Confidence ratings are my judgement, not anyone's published estimate.

Bet	What it rests on today	Stated target	Confidence	If it misses
Optical switching in the thalamus router	Mature in data centres already	Shipping now	High	Nothing: this is effectively a 2026 part
Chip-scale optical clock	Active commercialisation	Industry roadmaps	High	Keep the crystal oscillator; the loss is negligible
Entorhinal theta inside the assembloid	Assembloids already generate theta-gamma coupling	No date; the result exists	High	Pace theta electrically, as in the 2026 build
Photonic inference at useful scale	A fully integrated photonic processor doing every key deep-network computation , and programmable nonlinearity on chip	None public; demos are tiny (4 inputs, 3 outputs)	Medium	Keep GPUs. Costs power and latency, not function
Analog in-memory crossbar for the cerebellum	Long-running work at IBM and others	No single public date	Medium	FPGA filters already hit 1 kHz; this was a power win
Photo-uncaging with millisecond timing	Established chemistry, routine in slice work	No date	Medium	Microfluidic delivery: slower, still adequate
Perfused, vascularised assembloid running for years	Endothelial networks with blood-brain-barrier features and up to 3x lower cell death ; perfusion chips in progress	Projects running since 2025; no delivery date	Medium-low	The electronic twin and staggered cultures already cover this. The unit stays a consumable

Bet	What it rests on today	Stated target	Confidence	If it misses
Phase-change photonic synapses at scale	Optical synapses with Hebbian and STDP learning , about 70 per chip	No date	Low	In-place learning stays on GPUs. This is the weakest optical bet
Fault-tolerant quantum co-processor	Nord Quantique targets fault tolerance in 2030 ; Xanadu has published a blueprint for 2028-2029	2028-2031, depending on the company	Low-medium	Keep the annealer. It is planner-only either way, so nothing breaks
Diamond archive at practical cost	1.85 TB/cm3, millions-of-years retention , but not commercially viable on equipment cost	None	Low on cost, high on function	Use SSDs. You lose permanence, not capacity

The pattern worth noticing. Every low-confidence bet has a working fallback that costs power, money or elegance, but not capability. That is not luck; it is the point of the interface contract. The one bet where a miss genuinely hurts is the vascularised assembloid, because a memory unit that must be replaced twice a year is a different product from one that runs for years.

One bet deliberately left out. Running cortical function on living tissue at scale. There is no published result that makes that a 2030 proposition, and including it would make the whole column untrustworthy.

Whole-system comparison

All figures are design estimates for budgeting, not quotes. The body is the same small humanoid in both columns; only the brain rack changes.

	2026 build	2030 build
On-body electronics	115–205 W	70–120 W

	2026 build	2030 build
Brain rack compute	120–300 W	40–100 W
Wet module life support	200–400 W	200–400 W
Imaging and optical write	150–250 W	100–150 W
Total continuous power	585–1,155 W	410–770 W
Brain rack footprint	2 full racks	1 rack
Hardware cost	\$1.5M–\$3M	\$1M–\$2M, plus whatever the quantum access costs
Memory unit service interval	Every 3–6 months	Target: years
Perception latency	50–300 ms	10–50 ms
Planning cycle	1–10 s	0.2–2 s
Reflex and balance latency	Under 2 ms	Under 2 ms

Power barely halves, and the wet modules are why. Life support is the largest single line in both columns and does not improve, because keeping tissue alive at 37 degrees with continuous perfusion is a thermodynamic cost, not an engineering inefficiency. A fully dry build would draw far less; the wet memory unit is bought at that price.

The reflex row does not change, and that is the design working. The fastest loop is already as fast as it needs to be in 2026. Every 2030 improvement lands in the slow layers, where the headroom actually is.

For scale: the human brain runs on about 20 W. The 2030 build is still roughly 20–40 times that, for far less capability.

Still unsolved in 2030

Assume every bet above lands. Four things are still open, and they are the ones worth planning research around rather than purchasing.

1. Continual learning without forgetting. The 2030 build learns faster in more places, but nothing in it solves the problem of absorbing a new skill without degrading an old one. The staggered-culture and electronic-twin scheme works around it; it does not fix it. This

remains the single hardest problem in the whole design, and it is a software problem, not a parts problem.

2. Whether the tissue is doing anything silicon could not. The wet memory unit is justified by one-shot learning. Nobody has yet shown an organoid storing and recalling an episodic memory the way an adult hippocampus does — current hippocampal organoids most closely resemble fetal tissue. If that result does not arrive, Variant D is a very expensive way to get a fast-write database, and the honest move is to drop the wet modules.

3. Tuning the state bus. Four scalars standing in for dopamine, noradrenaline, serotonin and acetylcholine are trivial to wire and very hard to set. There is no principled method for choosing their values, and a bad setting produces a robot that is either reckless or frozen. No 2030 part helps here.

4. Verification of a system that rewrites itself. The safety architecture is sound for fixed software. A brain whose cortex updates its own weights while running, and whose memory unit is living tissue, cannot be certified by the methods that certify the brainstem supervisor. This is why the override rule matters more in D than in any earlier variant: the only verifiable claim is that the fast, fixed layers can always stop the slow, learning ones.

The emulation question is not on this list, because it is not close. A projection based on trends in supercomputing, connectomics and recording puts [mouse whole-brain simulation at cellular level around 2034, marmoset around 2044, and human later still](#), and a whole mouse connectome is described as [a 10–15 year goal](#). Variant D is a functional brain, not a copy of one, and that distinction holds through 2030 and well past it.