BLOOD: A Biological Root of Trust and Dynamic Memory Architecture for Advanced Computing Platforms

Milan Cheeks SELF Labs Inc ORCID: 0009-0008-7507-0050 Corresponding author: milan@selflabs.xyz

Preprint — architecture and position paper. Not yet peer reviewed. The system described is patent pending (U.S. Patent Application No. 18/901,437).

Abstract

Every recovery mechanism in modern computing is itself software or firmware, and therefore lives inside the same attack surface it is meant to defend. A robot or advanced computing platform whose operating system, models, or update channel are compromised has no path back to a known-good state that the compromise cannot also reach: the reset button is drawn in the same medium as the failure. This paper presents BLOOD (Biological and Lightweight Organics Optimized for Dynamic memory systems), a bio-nanotechnology architecture under development at SELF Labs Inc that integrates engineered exosomes with optogenetic control to form a vein-like biological network inside computing platforms and humanoid robots. BLOOD is proposed for two coupled functions. First, as a dynamic memory and signaling substrate: exosome-mediated molecular communication and light-addressed control provide memory storage, inter-component communication, and adaptive processing elements that complement, rather than replace, conventional electronic architecture. Second, and as this paper's central thesis, as a biological root of trust: because the biological layer sits physically outside the digital attack surface, it can hold a known-good recovery state and force a compromised platform back to that state through a channel that malware, by construction, cannot traverse. We describe the architecture at the level of its disclosed elements, analyze the security argument and its limits, survey the intended application classes (humanoid robotics, advanced computing, and computational-neuroscience research platforms), and specify a staged validation program. BLOOD is at the architecture stage: no bench results are reported here, and the paper distinguishes throughout between what the design asserts, what the underlying biological literature supports, and what remains to be demonstrated.

Keywords: biological computing; exosomes; optogenetics; root of trust; assured autonomy; humanoid robotics; molecular communication; neuromorphic systems

1. Introduction: the recovery problem

Autonomous systems are acquiring the ability to act in the physical world faster than our ability to recover them when they go wrong. A humanoid robot or an advanced AI platform can fail in two distinct ways. It can fail mechanically or electrically, which engineering handles well. Or it can fail *in its mind*: a compromised update, a poisoned model, a manipulated planner, or a hallucinating controller producing confident, dangerous behavior. The second failure class has a property the first does not. Every tool available to correct it — the watchdog, the rollback image, the remote management channel, the factory reset — is implemented in the same digital substrate that is compromised. Whoever or whatever controls the software controls the recovery from the software.

The trusted-computing answer to this problem is the hardware root of trust: a small, hardened component (a TPM or secure enclave) that anchors measurement and recovery [1]. Hardware roots of trust are a genuine advance, but they remain electronic components on the same boards, reachable through the same buses, firmware supply chains, and physical interfaces as everything else, and the history of enclave and TPM attacks shows that "separate chip" is a weaker separation than "separate physics."

BLOOD starts from the stronger separation. Its premise is that a recovery and memory layer implemented in *biology* — in engineered extracellular vesicles and light-controlled cellular machinery, physically present inside the platform but not addressable by any instruction the platform's processors can execute — is unreachable by digital compromise in principle, not merely in practice. Software cannot patch a vesicle. A malicious model cannot emit the optical and biochemical sequence that speaks to the biological layer unless the platform's legitimate control path grants it, and that path can be made narrow, physical, and auditable.

This paper makes three contributions. First, it states the BLOOD architecture as disclosed by SELF Labs: engineered exosomes and optogenetics integrated into a vein-like network providing memory, communication, and processing functions for advanced computing platforms and humanoid robots (Sections 3–4). Second, it develops the root-of-trust thesis rigorously, including the threat model under which the biological layer helps and the attack classes it does not address (Section 5). Third, it is explicit about evidence: BLOOD is an architecture with a patent application, not a demonstrated system, and the paper closes with the staged experimental program that would be required to change that (Sections 8–9).

2. Background

Exosomes and extracellular vesicles. Exosomes are nanoscale extracellular vesicles, released by cells, that carry proteins, lipids, and nucleic acids between cells and can transfer functional RNA cargo to recipient cells [2][3]. They are, in engineering terms, a manufactured-by-biology packet network: addressed by surface composition, loaded with molecular payloads, stable in circulation, and capable of crossing barriers (including, relevantly for neuroscience applications, the blood-brain barrier) that synthetic carriers cross poorly. The field's characterization standards are consolidated in the MISEV guidelines [3]. Engineered vesicles with designed cargo and targeting are an active research area in drug delivery; BLOOD's proposal is to use the same substrate as an information and state medium inside machines.

Optogenetics. Optogenetics gives genetically encoded, light-activated control of cellular processes with millisecond temporal precision and cellular spatial resolution [4][5]. For a computing architecture, its significance is the interface: light is a control channel that crosses the boundary between electronics and biology without wires, electrodes, or shared protocols — an optical air gap that can still carry structured commands.

Biological and molecular computing. Computing with molecules has a long lineage, from DNA-based computation [6] to molecular automata that evaluate state inside cells [7]. BLOOD differs from that lineage in intent: it is not proposed as a general-purpose replacement for electronic computation, where biology's speed disadvantage is decisive, but as a specialized substrate for the two functions where its properties are advantages rather than compromises — persistent state held outside the electronic system, and control exercised across a physics boundary.

Roots of trust and assured autonomy. The recovery framing draws on trusted computing [1] and on the assured-autonomy literature's concern with systems that must remain safe under compromise and distribution shift. BLOOD's claim is narrow and architectural: it proposes moving one specific function — custody of a known-good state and the authority to force a return to it — into a substrate the digital system cannot write to.

3. Architecture overview

As disclosed, BLOOD integrates three elements inside a host platform [8]:

1. **A vein-like biological network.** A bio-inspired network distributed through the computing platform, analogous in layout to vasculature: a transport and communication fabric reaching the subsystems it serves, rather than a single centralized component.
2. **Engineered exosome signaling.** Vesicles serve as the network's mobile elements, carrying molecular signals and state between nodes — the substrate for memory storage, inter-component communication, and adaptive processing behavior.
3. **Optogenetic control.** Light-addressed control points give the host platform's legitimate control path a structured way to write to, read from, and direct the biological layer, with the precision optogenetic techniques provide.

The host relationship is deliberately asymmetric. The electronic system is the fast, general computer; the biological layer is slow, specialized, and authoritative over a small set of functions. BLOOD is not proposed to run the robot's planner, perception, or control loops. It is proposed to (a) hold dynamic memory and signaling functions where biological persistence and parallelism are useful, and (b) hold the recovery authority described in Section 5. The public description of BLOOD additionally identifies enhanced connectivity, plasticity, and synchronization for AI systems as design objectives of the architecture [8]; consistent with the evidence position in Section 8, this paper treats those as objectives the validation program must test, not as properties already demonstrated.

The public record does not disclose, and this paper does not invent, the specific materials, vesicle engineering methods, optical parameters, network topology, or node design. Those elements are the subject of the patent application and of the staged program in Section 9.

4. BLOOD as a dynamic memory system

The name encodes the first function: Biological and Lightweight Organics Optimized for Dynamic memory systems. Three properties motivate a biological memory substrate as a complement to electronic memory:

- **Persistence without power or addressability.** Molecular state persists without refresh power and, critically for Section 5, without being mapped into the electronic system's address space. State the processor cannot

address, it cannot be tricked or compromised into rewriting.

- **Massively parallel molecular signaling.** Vesicle populations act in parallel by nature. The design objective is memory and signaling behavior — association, decay, reinforcement — implemented by population dynamics rather than by clocked access, a neuromorphic rather than von Neumann style of storage.
- **Plasticity as a native behavior.** Biological networks modify their own connectivity and response over time. For AI systems, a substrate whose memory physically remodels under use is a candidate medium for long-horizon adaptation that electronic systems currently simulate at energy cost.

The honest limits belong in the same section. Biological timescales are slow against electronic ones; molecular signaling will not compete with DRAM or cache on latency, and no serious version of this architecture proposes that it should. The claim is complementarity: electronic memory for speed, biological memory for persistence, parallelism of a different kind, and isolation.

5. BLOOD as a biological root of trust

This is the function SELF Labs regards as the architecture's center, defined in the company's framing work and stated publicly here: **BLOOD acts as a hack-proof factory reset for robots and autonomous platforms.**

Threat model. The adversary controls some or all of the platform's software stack: a malicious or corrupted update, a compromised model or planner, or remote exploitation of the operating system. The adversary does *not* have physical custody of the platform, and cannot synthesize biological materials or operate laboratory equipment against it. This threat model — remote, digital, total over software — is the standard one for fleet robots and connected autonomous systems, and it is the one under which the argument below holds. Section 5.3 states where it fails.

The recovery protocol, architecturally. The biological layer holds, in molecular state, the platform's known-good recovery configuration: the minimal trusted state from which the electronic system can be re-imaged and re-attested. A reset proceeds as follows. (i) A reset is *triggered* through a path that does not depend on the compromised software: a physical control, an out-of-band command authenticated at the optical interface, or an autonomous trigger from the biological layer's own state (for example, failure of a scheduled biological check-in that healthy host software maintains). (ii) The biological layer *asserts recovery authority*: through its interface to the host's power and boot chain, the platform is forced into its minimal recovery environment — the one action the layer must be able to compel electronically, implemented as a small, fixed, non-programmable interface. (iii) The known-good state held biologically is used to *verify and restore* the electronic system, and the platform reboots into attested software or remains inert. At no step does compromised software get a vote: the trigger path, the authority assertion, and the custody of the good state all sit outside its reach.

Why biology, specifically. Three properties do the work. *Unreachability:* there is no instruction set, driver, or network protocol by which software addresses vesicle state; the write path is optical and biochemical, gated by the physical interface. *Non-clonability of state:* the recovery state is held in a physical molecular configuration, not in a file that exfiltrated credentials can reproduce elsewhere. *Independent failure physics:* the layer does not share failure modes with the electronics — it does not crash when the OS crashes, is not patched when the fleet is patched, and cannot be rolled out to an attacker by a compromised update server, because it receives no updates.

5.3 Limits of the argument

The claim "hack-proof" is scoped, and the scoping matters. The biological root of trust does not help against an adversary with physical custody and laboratory capability, who can attack the biological layer on its own terms. It does not fix compromised *hardware* below the recovery interface, a malicious manufacturer, or a known-good state that was itself enrolled dishonestly — enrollment of the recovery configuration is a trusted ceremony and must be treated with the seriousness of key-signing. It introduces a new failure class of its own: biological components age, vary, and die, so the layer's own health monitoring and replacement lifecycle are first-order engineering problems, not afterthoughts. And the narrow electronic interface through which the layer compels recovery (step ii) is itself a digital component that must be kept minimal and fixed, because it is the one bridge the threat model cannot wish away. The paper's position is that a scoped, honestly bounded root of trust is worth more than an unscoped slogan — and that even under these limits, moving recovery authority across a physics boundary removes the largest and most common adversary (total remote software compromise) from the set of actors who can prevent a platform from being recovered.

6. Applications

The applications follow the two functions:

- **Humanoid robotics and autonomous platforms.** The primary case. Platforms that act physically, operate unattended, and receive over-the-air updates are exactly the systems exposed to unrecoverable software compromise, and exactly the systems where a forced return to a known-good state has safety value beyond its security value.
- **Advanced computing.** Dynamic biological memory and signaling as a complement to conventional architectures for complex simulations, large-scale data processing, and cryptographic workloads are stated objectives in the public description [8]; this paper records them as application targets whose feasibility the validation program, not this text, must establish.
- **Computational neuroscience and brain-computer interface research platforms.** The same elements — engineered vesicles as signaling and memory elements under optical control — constitute a research platform for neural-interface studies, the framing under which BLOOD fits biomedical research programs. For this application class the platform is a tool for studying biological information processing, not a computer peripheral, and it is the class where academic partnership is the natural execution path.

7. Safety and dual-use considerations

A technology that places engineered biological components inside machines carries obligations this paper states rather than defers. All biological work sits inside established biosafety practice: contained laboratory handling, characterized and non-replicating components as a design requirement for any deployed form, and institutional biosafety review for the experimental program. The root-of-trust function is defensive in design — its purpose is to return systems to a known-good state — but any technology that can force a platform's state deserves scrutiny for misuse in the other direction, and the enrollment and trigger authorities in Section 5 must be designed so that the power to reset cannot become a remote kill capability over someone else's machines: trigger authority should be physical or held by the platform's owner under the same custody rules as the enrollment ceremony. Finally, applications in neural-interface research inherit the ethics review structures of that field, and this paper's research-platform framing assumes them.

8. Current status and evidence position

Stated plainly, as a condition of taking the architecture seriously: BLOOD is patent pending (U.S. Application No. 18/901,437) and is offered by SELF Labs for qualified licensing evaluations [8]. The public record contains no bench data, no performance benchmarks, no named partners, and no demonstrated prototype, and SELF Labs' internal assessments describe the program as vision-stage, with the company's strength on the systems and computing side and wet-lab execution planned with qualified consultants, facilities, or an academic research partner. Nothing in Sections 3–6 should be read as a results claim. What exists today is an architecture, a patent application, a defined core function, and the validation program below. The underlying biological components — exosomes as information carriers, optogenetics as a control modality — are individually established in the literature [2][3][4][5]; the integration, and every system-level property claimed for it, is not yet established anywhere, including here.

9. Validation program

The program is staged so that each stage can falsify the next stage's premise cheaply:

1. **Signaling elements.** Characterize engineered vesicle populations as information carriers in vitro: cargo loading, targeting specificity, signal persistence, and read-out fidelity, against MISEV characterization standards.
2. **Optical addressing.** Demonstrate structured optogenetic control of the signaling elements: spatial and temporal selectivity, command repertoire, and error rates at the electronics-biology interface.
3. **Memory behavior.** Demonstrate the dynamic-memory claims at bench scale: state written optically, retained over defined intervals, read back reliably, and modified by population-level reinforcement and decay.
4. **Recovery prototype.** Integrate a minimal biological layer with a representative robotic subsystem and execute the Section 5 protocol end to end: compromise the subsystem's software, trigger recovery through the biological path, and verify return to the enrolled known-good state — including negative tests in which the compromised software actively attempts to block recovery.
5. **Security and lifecycle evaluation.** Independent red-team assessment under the stated threat model, plus the biological lifecycle program: component aging, variability across batches, health monitoring, and replacement.

Stages 1–3 are laboratory science suited to execution with an academic or contract research partner; Stages 4–5 are systems work in SELF Labs' own competence. The program is designed so that a failure at any stage produces

publishable knowledge about biological information substrates rather than a silent dead end.

10. Conclusion

BLOOD proposes that the most trustworthy part of a computing platform may be the part that is not a computer. As a dynamic memory system, it offers persistence, parallelism, and plasticity of a kind electronics simulates rather than has. As a biological root of trust, it addresses the recovery problem at the level where the problem actually lives: not in better software, but in a substrate software cannot reach. Both claims are, today, architectural. This paper has tried to state them at full strength and at honest scope in the same breath, because the staged program that could make them real begins with knowing exactly what has and has not been shown.

Declarations

Patent. The BLOOD architecture is the subject of U.S. Patent Application No. 18/901,437 (patent pending), held by SELF Labs Inc.

Competing interests. The author is the founder of SELF Labs Inc, which owns the BLOOD patent application and offers BLOOD for qualified licensing evaluations. This paper is a company architecture publication and should be read as one: its claims are the company's claims, stated under the evidence discipline of Section 8.

Data availability. No experimental data is reported in this paper; there is nothing to deposit. Bench data generated under the Section 9 program will be published with the corresponding results papers.

References

[1] Trusted Computing Group, "TPM 2.0 Library Specification," 2019.

[2] Valadi, H., et al., "Exosome-mediated transfer of mRNAs and microRNAs is a novel mechanism of genetic exchange between cells," *Nature Cell Biology*, 2007.

[3] Théry, C., et al., "Minimal information for studies of extracellular vesicles 2018 (MISEV2018)," *Journal of Extracellular Vesicles*, 2018.

[4] Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., Deisseroth, K., "Millisecond-timescale, genetically targeted optical control of neural activity," *Nature Neuroscience*, 2005.

[5] Deisseroth, K., "Optogenetics," *Nature Methods*, 2011.

[6] Adleman, L. M., "Molecular computation of solutions to combinatorial problems," *Science*, 1994.

[7] Benenson, Y., et al., "An autonomous molecular computer for logical control of gene expression," *Nature*, 2004.

[8] SELF Labs Inc, "BLOOD — Biological and Lightweight Organics Optimized for Dynamic memory systems," product and technology description, selflabs.xyz/blood, accessed November 19, 2026.