NEUROMORPHIC COMPUTING

Neuromorphic Silicon Deep Dive (Loihi 2 & SpiNNaker)

Current-source review, evidence-bounded adoption decision, explicit limitations, reproducibility controls, and preserved source research note.

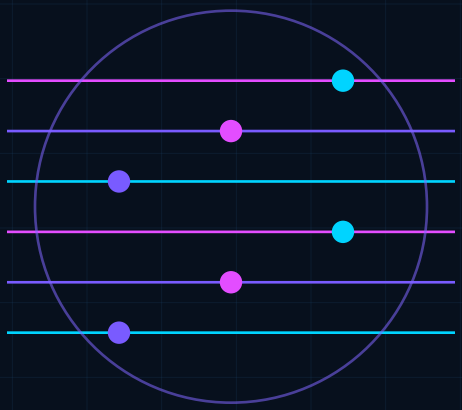

PERMANENT ID	EVIDENCE	ADOPTION	FRESHNESS
RES-017	B	MONITOR AND LAB VALIDATE	STABLE WATCH

Freshness review: 2026-09-17 / Next scheduled review: 2027-03-16

RESEARCH CONTROL SURFACE

Required Research Fields

Each field remains independently reviewable; evidence strength does not imply production authority.

EVIDENCE GRADE — **B**

Strength of the retained source set and technical basis.

SOURCE AUTHORITY — **2 registered authorities**

Intel Neuromorphic Computing Research; SpiNNaker2 Platform

FRESHNESS — **STABLE WATCH**

Reviewed 2026-09-17; dynamic sources require event-driven revalidation.

CONFLICTS — **VISIBLE / NOT SILENT**

Differences between specification, vendor behavior, research claims, and reproduced evidence must remain explicit.

LIMITATIONS — **EXPLICIT**

No certification, procurement approval, legal conclusion, or unbounded production claim is created by this report.

REPRODUCIBILITY — **REQUIRED**

Material claims require exact versions, representative data, failure cases, artifacts, and repeatable acceptance criteria.

ADOPTION POSTURE — **MONITOR AND LAB VALIDATE**

Maintain neuromorphic systems as a research track for event-driven sensing and ultra-low-power temporal workloads. Require hardware access, toolchain validation, and energy-normalized benchmarks.

REVIEW TRIGGER — **2027-03-16**

Scheduled at 180 days; earlier on material source, vulnerability, implementation, or contradictory-evidence change.

Authority Register and Freshness Delta

Primary-source lineage is preserved; current-source changes are separated from unperformed implementation claims.

AUTHORITY 01

Intel Neuromorphic Computing Research

https://www.intel.com/content/www/us/en/research/neuromorphic-com

AUTHORITY 02

SpiNNaker2 Platform

https://spinncloud.com/spinnaker2

2026-09-17 FRESHNESS DELTA

Primary-source register retained and freshness controls advanced to the current review date. No new product or laboratory performance claim is introduced without reproduced evidence.

Freshness policy: scheduled review + immediate review on source supersession, material release, vulnerability, retraction, or contradictory evidence.

Evidence Must Survive Contact With Reality

- Source authority is not implementation evidence. Product behavior, scale, security, interoperability, and operator fit remain environment-specific.
- Conflicting evidence is recorded rather than averaged away. Specification behavior, vendor implementation, preprint results, and production observations are separate evidence classes.
- Quantitative claims are informative until reproduced under documented hardware, software, data, topology, load, and failure conditions.
- Adoption posture is reversible. A stronger source set can coexist with a pilot-only or research-only recommendation when maturity, safety, economics, or operability remain unresolved.

ADOPTION / VALIDATION / REVIEW

Decision Gate and Validation Plan

ADOPTION POSTURE

MONITOR AND LAB VALIDATE

Maintain neuromorphic systems as a research track for event-driven sensing and ultra-low-power temporal workloads. Require hardware access, toolchain validation, and energy-normalized benchmarks.

1 / SOURCE	Pin exact versions, publication state, retrieved artifacts, and source hashes.
2 / PROTOTYPE	Demonstrate the core mechanism with representative data and instrumented execution.
3 / FAILURE	Exercise malformed input, dependency loss, stale state, overload, recovery, rollback, and misuse.
4 / BASELINE	Compare against an incumbent, classical, or simpler architecture using the same workload.
5 / ACCEPT	Record acceptance criteria, residual limitations, owner, expiration, and revalidation triggers.

EXPIRATION / REVIEW TRIGGER

Next scheduled review: 2027-03-16 (180-day cadence). Review sooner after material source revision, breaking implementation release, critical vulnerability, retraction, benchmark contradiction, changed workload, or changed legal/safety boundary.

8. Complete Source Research Note

SOURCE-LINEAGE NOTICE

The following material preserves the complete original source note, excluding only the conversational wrapper and outer Markdown fence. Legacy status labels and absolute language are retained for traceability and do not override the candidate status, limitations, or adoption decision in this edition.

Document ID: RES-017 Version: 1.0.0 (Definitive Registry Edition) Status: Published Research Note Classification: Public Organization: Mythos Systems (MYTHOS AI) Owner: Aaron Stovall Effective Date: 2026-07-16 Applies To: Edge AI engineers, embedded systems architects, and robotics specialists designing continuous-learning, ultra-low-power physical AI, and sensor-fusion systems Companion Standards: MYTHOS-EMT-0003 (Neuromorphic Computing & NPU Standard), MYTHOS-EMT-0002 (Sovereign AI), MYTHOS-IOT-0001 (IoT & Edge Sensors), MYTHOS-ROB-0001 (Physical AI)

The architecture of edge AI has failed. Organizations attempt to run continuous, always-on inference—such as visual tracking, audio wake-word detection, and robotic sensor fusion—using Graphics Processing Units (GPUs). These GPUs are synchronous, dense matrix multipliers that consume hundreds of watts of power and generate massive heat. Brute-forcing continuous edge inference with GPUs is architecturally unsustainable and economically unviable for battery-operated or deeply embedded sovereign systems.

Neuromorphic silicon, spearheaded by Intel Loihi 2 and SpiNNaker 2, shatters this paradigm. Inspired by the biological brain, these chips are fully asynchronous, event-driven Spiking Neural Network (SNN) accelerators. They do not process data in synchronized frames; they process discrete "spikes" in real-time. Because power is only consumed when a spike occurs, neuromorphic chips achieve sub-milliwatt continuous inference.

This research note defines the architectural dynamics of neuromorphic silicon. It dissects the mechanics of Spike-Timing-Dependent Plasticity (STDP) for local online learning, the routing constraints of Network-on-Chip (NoC) architectures, and the direct mapping of event-based sensors (Dynamic Vision Sensors) to silicon neurons. Understanding these dynamics is a prerequisite for building the ultra-efficient, adaptive physical AI mandated by the Mythos Emerging Tech Standards.

To understand neuromorphic compute, we must move from dense, synchronous matrix math to sparse, asynchronous temporal dynamics.

1.1 Spiking Neural Networks (SNNs)

Traditional ANNs (Transformers, CNNs) use continuous, floating-point activation values (e.g., `0.87`). SNNs use discrete, binary temporal events called "spikes."

- Leaky Integrate-and-Fire (LIF) Neurons: A neuromorphic neuron accumulates incoming spikes over time. When its membrane potential crosses a threshold, it fires a spike to downstream neurons, and its potential resets.
- The Paradigm Shift: Information is not encoded in the amplitude of the activation, but in the timing and frequency of the spikes.

1.2 Intel Loihi 2 (The Digital Microcode Brain)

Loihi 2 represents the "digital neuromorphic" paradigm. It utilizes a massively parallel mesh of cores, each hosting thousands of LIF neurons.

- The Architecture: It features a Network-on-Chip (NoC) that routes spikes asynchronously between cores.
- The Advantage: Loihi 2 allows for programmable microcode learning rules. Engineers are not locked into a static model; they can define custom local learning algorithms (like STDP) that update synaptic weights directly on the silicon in real-time.

1.3 SpiNNaker 2 (The Massively Parallel ARM Brain)

SpiNNaker (Spiking Neural Network Architecture) represents the "biologically plausible" paradigm. It is built from thousands of embedded ARM cores specifically designed to simulate biological neural networks in real-time.

- The Architecture: Unlike Loihi's custom digital neurons, SpiNNaker uses general-purpose ARM cores running software simulations of neurons, communicating via a specialized packet-switched fabric.
- The Advantage: Extreme flexibility. Any biological neuron model (e.g., Hodgkin-Huxley) can be programmed and simulated in real-time, making it ideal for computational neuroscience and highly complex, dynamic AI research.

While neuromorphic chips offer sub-milliwatt efficiency, they introduce severe physical and algorithmic constraints that standard GPU developers do not face.

2.1 The Non-Differentiable Training Boundary

Standard deep learning relies on backpropagation, which requires smooth, differentiable mathematical functions. Spikes are discrete (0 or 1), non-differentiable step functions.

- The Flaw: You cannot natively backpropagate an error gradient through a spiking neuron. Training SNNs via standard PyTorch/TensorFlow is mathematically broken.
- The Mitigation: SNNs MUST be trained using either Surrogate Gradient methods (approximating the spike derivative during training) or biologically local learning rules like STDP, which do not require global error propagation.

2.2 The Network-on-Chip (NoC) Routing Bottleneck

Because neuromorphic chips host millions of neurons on a single die, spikes must be routed across a physical mesh. If millions of neurons fire simultaneously, the NoC experiences routing congestion.

- The Flaw: If a spike is delayed due to network congestion, the temporal information is corrupted. The SNN's accuracy degrades.
- The Mitigation: The architecture MUST utilize topology-aware mapping. Neurons that heavily communicate (e.g., layers of a vision pipeline) MUST be physically placed on adjacent cores to minimize NoC hop count.

2.3 The Tooling Immaturity

The software ecosystem for GPUs (CUDA, cuDNN, vLLM) is mature and highly optimized. The neuromorphic ecosystem (e.g., NengoDl, Lava) is nascent.

- The Constraint: Compiling a standard ANNs (like ResNet) into an SNN that runs efficiently on Loihi 2 requires specialized knowledge of spike-encoding and neuron parameter tuning. It is not a "one-click" deployment.

The true power of neuromorphic silicon is realized when paired with event-based sensors and local learning rules.

3.1 Spike-Timing-Dependent Plasticity (STDP)

In a standard ANN, learning requires pausing inference, computing gradients, and updating weights on a GPU. In a neuromorphic chip, learning happens in situ via STDP.

- The Mechanic: If a pre-synaptic neuron fires just before a post-synaptic neuron, the synapse between them is strengthened (Long-Term Potentiation). If it fires just after, the synapse is weakened (Long-Term Depression).
- The Impact: The chip learns continuously, in real-time, without a global loss function or backpropagation. A robot can adapt to a new environment dynamically, meeting the continual learning mandates of `MYTHOS-AI-0012`.

3.2 Direct Sensor-to-Spike Mapping

Standard cameras output synchronized frames (e.g., 60 FPS). This forces the AI to process redundant pixels of a static background.

- Event-Based Sensors: Dynamic Vision Sensors (DVS) only fire when a pixel's log-intensity changes. They do not output frames; they output asynchronous spikes.
- The Synergy: A DVS sensor can be directly wired (logically) to a Loihi 2 chip. A spike from the sensor triggers a spike in the silicon neuron with zero CPU overhead and zero frame-buffering latency. This enables microsecond reaction times for high-speed robotics and physical AI (`MYTHOS-ROB-0001`).

The fundamental shift of neuromorphic compute is the decoupling of intelligence from power consumption.

- GPU Edge Inference: A Jetson Orin GPU consuming 30W can run complex vision models, but it drains a drone battery in 20 minutes and requires active cooling.
- Neuromorphic Inference: A Loihi 2 chip consuming <1 milliwatt can run an SNN that tracks objects and adapts to lighting changes in real-time. It drains a drone battery in days, not minutes, and requires no cooling.

The Scaling Law: As neuromorphic models scale to millions of neurons, the power consumption scales sub-linearly. Because 90% of neurons are silent at any given moment (sparse activation), adding more neurons only marginally increases power draw. This allows organizations to deploy highly intelligent, always-on AI into micro-IoT sensors, tactical edge devices, and autonomous swarms, achieving the Mythos mandate for sovereign, ultra-low-power edge AI.

To deploy neuromorphic silicon in production, the Mythos Architecture enforces strict hardware-aware routing and hybrid execution models.

5.1 Hybrid Neuromorphic-LLM Routing

SNNs excel at continuous, low-power perception (e.g., listening for a wake word, tracking a moving object), but they lack the dense, in-context reasoning of LLMs.

- The Mitigation: The edge device MUST utilize a heterogeneous architecture. The neuromorphic chip runs continuously, sub-milliwatt, filtering noise and parsing raw sensor data into semantic events. Only when a high-level cognitive trigger occurs (e.g., "Intruder detected") does the device power on the GPU/LLM to perform complex reasoning and generate an alert. This combines the efficiency of SNNs with the intelligence of Transformers.

5.2 Topology-Aware SNN Compilation

The Mythos compilation pipeline MUST treat the physical layout of the neuromorphic chip as an architectural invariant.

- The Mitigation: Compilers (like Intel Lava) MUST analyze the SNN graph and map highly connected neuron clusters to the same physical core or adjacent cores. Cross-die spike routing MUST be minimized to prevent NoC congestion and temporal latency.

5.3 Local Continual Learning Bounds

Because STDP allows the chip to learn continuously, an SNN is susceptible to drift or adversarial manipulation of its environment.

- The Mitigation: The architecture MUST bound the STDP learning rate and synapse weight bounds. Furthermore, the system MUST periodically snapshot the SNN's synaptic state to non-volatile memory, allowing rollback if the chip learns a malicious or hallucinatory pattern.

Synchronous, frame-based AI compute is a brute-force artifact of the GPU era. For continuous, edge-native, and physical AI, it is an architectural dead end. Neuromorphic silicon, via Spiking Neural Networks and local learning rules like STDP, brings compute back to the laws of biology. By processing data as sparse, asynchronous events, we achieve sub-milliwatt

intelligence that adapts in real-time, unlocking the next generation of sovereign, deeply embedded autonomous systems.

```
A GPU computes frames; a neuromorphic chip computes reality.
A global gradient is a bottleneck; local STDP is a flow.
A synchronous clock is a power tax; an asynchronous spike is free.

The synapse is not a static weight; it is a fluid function of time.

> Perception may be continuous.
> Edge efficiency must be absolute.
```

End of Research Note RES-017

9. Source Register

Source	Authority and Status	Freshness Control
Intel Neuromorphic Computing Research https://www.intel.com/content/www/us/en/research/neuromorphic-computing.html	Primary vendor research Living research program Published: Living	Validated: 2026-07-22 Review on material revision, withdrawal, supersession, or scheduled report review.
SpiNNaker2 Platform https://spinncloud.com/spinnaker2	Primary project/vendor documentation Living Published: Living	Validated: 2026-07-22 Review on material revision, withdrawal, supersession, or scheduled report review.

10. Lineage, Review, and Publication Record

Record	Value
Original source file	RES-017_Neuromorphic_Silicon_Deep_Dive_Loihi_SpiNNaker.md
Original source words	1564
Upgrade type	Enterprise research report; source and freshness validation; limitations; adoption decision; reproducibility plan
Generated on	2026-07-22
Current status	Candidate Research Report
Publication authority	Formal Mythos approval not yet recorded
Next review	180 days or event-driven trigger, whichever occurs first