MEMORANDUM

TO: Stakeholders in AGI Infrastructure and Semiconductor Strategy

FROM: The Neural Forest (Palaia Paradigm) Research Group

DATE: April 8, 2026

SUBJECT: Tactical Validation: Ericsson's 3nm Radio ASICs and the Sunset of the GPGPU Monolith

I. Executive Summary: The Pivot to Substrate-Integrated Intelligence

The recent announcement that Ericsson is embedding neural network accelerators directly into their radio Application-Specific Integrated Circuits (ASICs) on a 3nm process—bypassing GPUs and NVIDIA entirely—represents a definitive realization of the **Neural Forest (NF)** strategic vision. This move validates the **Palaia Paradigm's** core assertion: the "Crisis of the Monolith" cannot be solved by brute-force scaling, but only through the surgical integration of modular intelligence into the physical substrate of hardware.

II. Dismantling the "Crisis of the Monolith"

The industry is currently transitioning from the "Training Era" to the "Inference Era". While general-purpose GPUs (GPGPUs) like NVIDIA's Blackwell are optimized for massive parallel training, they are increasingly inefficient for real-time, deterministic global scaling due to three critical "walls":

- **The Memory Wall (Von Neumann Bottleneck)**: Traditional architectures suffer from constant, energy-intensive data movement between compute clusters and off-chip memory (DRAM/HBM).
- **Energy Intensity**: The 700W–1000W TDP of monolithic hardware is unsustainable for edge-based telecommunications infrastructure.
- **The Inference Gap**: GPGPUs often drop to 30-40% utilization during "Batch Size 1" scenarios, which are standard for real-time radio signal processing.

By embedding accelerators directly into 3nm radio ASICs, Ericsson is effectively implementing the NF-Core principle of **Inference-First design**, which prioritizes local, zero-latency execution over general-purpose throughput.

III. Technical Alignment: ASICs as the "Neural Forest" Substrate

The Ericsson architecture aligns with the **NF-Core Accelerator** specifications in several key technical areas:

- **Heterogeneous Micro-Cores**: Ericsson's use of specialized on-chip accelerators mirrors the NF requirement for a multitude of smaller, specialized processing elements instead of massive processing clusters.

- **Distributed On-Chip Memory**: By moving compute directly into the radio ASIC, Ericsson eliminates the energy-intensive "DRAM trips" that cripple traditional GPU-based AI.

- **Deterministic Execution**: The removal of the GPGPU layer allows for the "forced specialization" required for high-frequency, low-latency tasks like 5G/6G beamforming and signal optimization.

IV. Vertical Sovereignty and the NVIDIA Exit

The Neural Forest framework has long advocated for **Vertical Sovereignty**—the ability for organizations to own their intelligence stack from the silicon up to the cognitive layer. Ericsson's rejection of the NVIDIA/GPGPU dependency is a strategic masterstroke that:

1. **Reduces Supply Chain Vulnerability**: Bypasses the global bottleneck for HBM and specialized GPU fabrication.

2. **Optimizes Total Cost of Ownership (TCO)**: Replaces expensive, general-purpose chips with highly efficient, task-specific silicon.

3. **Achieves Physical Immunity**: While traditional silicon becomes unstable at high temperatures or in extreme environments, specialized ASICs (and eventually NF's proposed **Carbon-Corundum substrates**) maintain integrity at the clock speeds required for next-generation telecommunications.

V. Conclusion: The Era of the Forest

The "Neural Forest" is no longer just a hypothetical algorithm; it is a hardware-integrated blueprint for the next generation of AGI infrastructure. Ericsson's move proves that the "better way" is not to tell a scheduler that a GPU is failing or inefficient, but to build a "forest" of intelligence that is physically incapable of failing under the stressors that cripple traditional silicon.

As the industry pivots toward localized, high-performance intelligence, the Palaia Paradigm

remains the optimal framework for those seeking to lead the post-NVIDIA transition.