

The Physics of DRAM: A Leaky Bucket

1T1C Architecture

1 Transistor + 1 Capacitor.
Incredibly simple, allowing
massive density (32GB+ per stick).
But the capacitor holds less than
0.7 femtofarads (fF) of charge.
It constantly drains.




The Refresh Tax

To prevent the memory from dying and turning a 1 into a 0, the chip must read and rewrite every single cell every 32 to 64 milliseconds.

This brute-force regeneration burns roughly 10% of the total DRAM power budget just to keep the data alive.

Volatility: Forgetting in a Fraction of a Second



Power ON

Continuous Refresh

Constant rewriting means DRAM is extremely **power-hungry**, eating up energy even when the chip is idle.

Power OFF

Instant Volatility

Cut the power, and the refresh tax stops. The charge drains entirely in well under a second. Your RAM forgets instantly. Speed and **volatility** are two sides of the exact same physical coin.

The 2025 Market Inversion: Old Tech Outprices the New



DDR5 & HBM (The New Tech)

Status: AI Market Focus

Dynamic: Swallowed global fab capacity. High production volume by Big Three (Samsung, SK Hynix, Micron).

DDR4 (The Dead Tech)

Status: Last-Time-Buy Panic


Dynamic: March 2025: CXMT pivots away from DDR4, yanking cheap supply. Buyers hoarding for legacy servers and industrial tech.

DDR4

July 2025 Price Inversion
+300% Price Jump

Spot prices for 16Gb DDR5 leaped from \$6.84 to \$27.20. By July, old DDR4 was actually selling for more than cutting-edge DDR5.

The Scaling Wall: Geometry, Not Lithography



The Escape Hatch: Stacking the Future

Step 1: Sideways Shrinking (Current)

The $6F^2$ layout is maxed out. SK Hynix's $4F^2$ Vertical Gate buys only a few more nodes. The 10nm wall is firm.



Step 2: Vertical Stacking (Early 2030s)

3D DRAM: Stop shrinking sideways, start stacking capacitors vertically (the flash memory playbook).

The Foundation of AI

High-Bandwidth Memory (HBM) isn't exotic new magic. It is just these same fragile, leaky DRAM dies, thinned out and stacked together by the thousands. Every trillion-dollar AI model is built on top of this exact same leaky bucket.