6. Green HPC & Energy-Efficient Computing

6.1 The Energy Challenge

Green HPC refers to the optimization of supercomputers and data centers to maximize energy efficiency, reduce carbon footprints, and minimize environmental impact.

Modern supercomputers consume enormous amounts of electricity. Frontier consumes approximately 21 megawatts — enough to power 20,000 homes. The global HPC market's energy footprint is a serious environmental and economic concern. Energy cost perspective At $0.10/kWh, a 10 MW data center costs $8.76 million/year just in electricity.

6.2 Key Metrics

Power Usage Effectiveness (PUE)

Power Usage Effectiveness, universally abbreviated as PUE, is the standard metric used across the entire data center and HPC industry to measure how efficiently a facility uses its electricity. The formula is straightforward:

PUE = Total Facility Power / IT Equipment Power

The logic behind this ratio is simple — a facility consumes power not just for its servers and compute nodes, but also for cooling systems, lighting, power distribution units, uninterruptible power supplies, and building infrastructure. All of that non-computing power consumption is overhead.
One significant achievement is the Google's data centers, which have invested heavily in custom cooling infrastructure and machine learning-driven energy management, operate at approximately 1.1 PUE, making them among the most energy-efficient large-scale computing facilities in the world.

Performance per Watt

While PUE measures facility efficiency, Performance per Watt measures the efficiency of the computing hardware itself — how much useful work you extract from every watt of power consumed by the machines.

Performance/Watt = FLOP/s / Watts

This is the defining metric of Green HPC. It shifts the conversation from raw speed — how many FLOP/s can a machine do — to sustainable speed — how many FLOP/s can a machine do per watt of power consumed.

6.3 Energy Efficiency Techniques

Hardware-Level Techniques

Energy efficiency in HPC begins at the hardware level, where processor architects and chip designers make decisions that determine the power envelope of the entire system.

- Dynamic Voltage and Frequency Scaling (DVFS) — reduce clock speed and voltage when full performance isn't needed
- Power gating: The ability to completely cut power to cores, memory controllers, or entire circuit blocks that are not currently active.
- ARM-based processors have long been known for exceptional performance-per-watt characteristics, originally developed for battery-powered mobile devices but now proving highly competitive in HPC.
- GPU and TPU accelerators deliver dramatically higher performance per watt than general-purpose CPUs for the specific mathematical workloads

System-level:
Even with the most efficient hardware, a poorly designed facility wastes enormous amounts of energy — and the biggest system-level energy consumer in any HPC center is cooling. Traditional air cooling blows chilled air through server racks, but air has poor thermal conductivity and requires powerful fans and compressors that themselves consume significant power.

- Liquid cooling instead of air cooling (more efficient, enables heat reuse)
- Immersion cooling (servers submerged in dielectric fluid)
- Free cooling using outdoor air when ambient temperature allows
- Waste heat recovery repurposing heat for building warming

Software-level:

Hardware and facilities set the boundaries of energy efficiency, but software determines how well the hardware is actually utilized within those boundaries —

and poor software choices can waste just as much energy as poor hardware choices. Power-aware scheduling — place workloads on energy-efficient nodes

- Job packing — avoid running machines at low utilization
- Algorithm selection — energy-efficient algorithms (fewer operations, better memory patterns)
- Checkpoint/restart to avoid recomputation

6.4 Exascale Energy Goals

The Exascale computing program set a clear, non-negotiable energy target from the beginning: achieve 1 ExaFLOP/s of sustained performance within a power budget of **20 to 30 megawatts**. This constraint was not arbitrary — it reflected both the practical limits of what a single facility's power infrastructure could realistically supply and the economic reality that electricity costs beyond this threshold would make operating an Exascale machine financially prohibitive for national laboratories. Frontier achieves this target. The broader aim is to push towards **100 GFLOP/s per watt**, which requires fundamental advances in chip architecture, memory technology, and interconnects. The path to Zettascale computing in the 2030s runs directly through this energy efficiency challenge — without fundamental breakthroughs in performance per watt, the power requirements of a Zettascale machine would be measured not in megawatts but in gigawatts, exceeding the output of large nuclear power plants and making such systems economically and environmentally indefensible.