

The QWAV Decade: Enterprise p-Adic Computing 2025-2035

Rowan Brad Quni-Gudzin

2026-07-31

Author: Rowan Brad Quni-Gudzin | **Date:** 2026-07-31 | **License:** QNFO-ULA: <https://legal.qnfo.org/>

Forward-Looking Statements: This document contains forward-looking statements based on current expectations, estimates, and projections. These statements are not guarantees of future performance and involve risks and uncertainties. Actual results may differ materially. The forecasts, timelines, and market assessments herein are based on publicly available data, published research, and calibrated subjective judgment as of 2026-07-31. All figures marked with a dagger (†) are design targets — not yet demonstrated in industry-wide practice.

Abstract

Computing efficiency — measured by Joules Per Computational Unit (JPCUB †) — will be the dominant competitive axis in enterprise computing by 2035. This document forecasts the JPCUB ecosystem's evolution across three eras. **Era 1: Benchmark (2025-2028)** — JPCUB is validated against public hardware data, and the first enterprise adopters deploy JPCUB in procurement. **Era 2: Adoption (2028-2032)** — JPCUB is cited in Requests for Proposal (RFPs), cloud providers publish JPCUB scores, and the first regulatory mandates appear. **Era 3: Dominance (2032-2035)** — JPCUB is a standard metric alongside Total Cost of Ownership, hardware is designed for JPCUB optimization, and p-adic computing primitives enter mainstream processor architectures. Each era maps onto concrete enterprise decision points: when to adopt JPCUB, when to optimize architecture for JPCUB, and when the competitive landscape shifts from performance to efficiency.

1. What Is JPCUB?

Joules Per Computational Unit (JPCUB †) is a hardware-independent, workload-normalized metric of computing efficiency. Unlike MFLOPS/Watt (Green500's metric, which favors architectures with high FLOP counts regardless of actual workload throughput), JPCUB measures energy per unit of computational work — where "computational unit" is defined by a standardized workload benchmark [CROSS-REF: jpcub-validation/paper.md].

1.1 Why JPCUB Now?

Three trends make JPCUB's timing critical:

- 1. Data Center Energy Growth:** Global data center electricity consumption is projected to reach 1,000 TWh/year by 2026 — approximately 3% of global electricity — and may double by 2030 [andreae2015global, doe2016usdatacenter]. Energy efficiency is no longer a "green" nice-to-have; it is an operational cost driver and a regulatory target.
- 1. Hardware Heterogeneity:** The era of "one CPU architecture for all workloads" is over. Enterprise computing spans CPUs (x86, ARM, RISC-V), GPUs (NVIDIA, AMD, Intel), TPUs (Google), FPGAs, and custom ASICs. A single efficiency metric that works across all of these is essential for procurement.
- 1. Green Computing Regulation:** The EU Energy Efficiency Directive, California's data center regulations, and similar frameworks globally are moving from "voluntary reporting" to "mandatory compliance." A

standardized efficiency metric will be required for compliance — and the metric that becomes the standard will shape the industry.

1.2 How JPCUB Differs from Existing Metrics

Metric	Measures	Limitation
MFLOPS/Watt (Green500)	Theoretical peak FLOPs per Watt	Favors architectures with high FLOP counts regardless of actual workload throughput
PUE (Power Usage Effectiveness)	Infrastructure efficiency (total facility power / IT equipment power)	Measures datacenter cooling/power distribution efficiency, not compute efficiency
SPECpower_ssj	Server-side Java workload efficiency	Single workload type; architecture-specific
TCO (Total Cost of Ownership)	Total cost including hardware, power, cooling, staffing	Aggregates efficiency into a dollar figure — loses the compute-normalized signal
JPCUB †	Joules per standardized computational unit, architecture-independent	Validated on public benchmark data; cross-architecture validation in progress

2. Three Eras of JPCUB Adoption

2.1 Era 1: Benchmark (2025-2028)

Key Developments		
Year	Milestone	Status
2025	JPCUB metric defined, initial validation on CPU benchmarks	✅ Complete [CROSS-REF: jpcub-validation]
2026	JPCUB cross-architecture validation (CPU vs GPU) published	In progress
2027	Open-source JPCUB benchmark suite released (GitHub + the QNFO archive) †	Planned
2028	≥1 enterprise IT department pilots JPCUB in procurement	Forecast
2028	JPCUB cited in ≥1 academic publication independent of QNFO	Forecast

Enterprise Adoption Profile — Era 1

Early adopters: Cloud providers and hyperscalers (AWS, Google, Microsoft) who already track energy efficiency internally. They will test JPCUB against their existing internal metrics. If JPCUB correlates well with their operational data, they will adopt it for internal use — but are unlikely to publish JPCUB scores publicly in this era.

Hardware vendors: Intel, AMD, and NVIDIA will monitor JPCUB. They will not adopt it in marketing materials in Era 1 (too early, no customer demand), but their internal benchmarking teams will evaluate it.

Enterprise IT: 1-3 Fortune 500 companies with strong sustainability mandates will pilot JPCUB in procurement. These are likely in financial services (driven by ESG reporting requirements) and technology (driven by technical capability).

Technology Stack — Era 1

- Benchmark software:** Open-source JPCUB suite (Python + compiled workloads: SPEC CPU, LINPACK, STREAM, MLPerf)

- **Data:** Public benchmark databases (SPEC, Geekbench, Phoronix) + cloud provider instance-type data
- **Validation methodology:** Cross-architecture normalization via workload-equivalence classes
- **Distribution:** GitHub + DOI + Jupyter notebooks for reproducible analysis

2.2 Era 2: Adoption (2028-2032)

Key Developments

Year	Milestone	Status
2029	JPCUB cited in ≥2 enterprise hardware procurement RFPs	Forecast
2030	AWS/GCP/Azure publish JPCUB scores for ≥50% of compute instance types	Forecast
2031	EU Energy Efficiency Directive amendment references JPCUB as a compliance metric	Forecast
2032	≥1 major chip vendor publishes JPCUB regression data for ≥1 product generation	Forecast

Enterprise Adoption Profile — Era 2

Cloud providers: AWS, GCP, and Azure publish JPCUB scores alongside instance types. Customers select instances by JPCUB-aware auto-scaling — migrating workloads to instances with the best JPCUB for the workload type. This creates competitive pressure: if AWS instances have better JPCUB than Azure for a workload class, Azure loses customers.

Hardware vendors: At least one major chip vendor (Intel, AMD, or NVIDIA) publishes JPCUB regression data showing that JPCUB improved generation-over-generation. This becomes a marketing advantage — "our chips are more EFFICIENT, not just faster."

Enterprise IT: JPCUB is a standard consideration in procurement RFPs alongside TCO. IT departments compare hardware and cloud vendors on JPCUB as the ENERGY dimension of procurement. TCO models include JPCUB as a line item: "energy cost per compute unit alongside dollar cost per compute unit."

Regulation: Government energy-efficiency standards for data centers adopt (or at minimum reference) JPCUB. New data centers must report JPCUB per workload class. This is the inflection point — regulation transforms JPCUB from a "nice-to-have" metric to a "must-have" compliance requirement.

Technology Stack — Era 2

- **Cloud integration:** JPCUB published via instance metadata APIs
- **Auto-scaling:** JPCUB-aware cloud orchestration (migrate workloads to instances with best JPCUB for workload type)
- **Regulatory compliance:** Standardized JPCUB reporting format (ISO or equivalent)
- **Hardware design:** JPCUB regression testing integrated into chip design flows alongside PPA (Power-Performance-Area)

2.3 Era 3: Dominance (2032-2035)

Key Developments

Year	Milestone	Status
2033	JPCUB is a standard metric alongside TCO in enterprise procurement	Aspirational
2033	≥3 major cloud providers publish JPCUB for 100% of compute instances	Aspirational
2034	p-adic arithmetic primitives (native p-adic number representation) enter ≥1 processor architecture	Aspirational
2035	JPCUB-optimized processor achieves 10× efficiency improvement over 2026 baseline at iso-workload †	Design target

Enterprise Adoption Profile — Era 3

Hardware vendors: Competition shifts from "more FLOPS" to "better JPCUB." A processor that is 10× more efficient at the same workload dominates the market. Chip architectures are designed for JPCUB from the ground up — p-adic number representations in ALU, hierarchical memory mapping onto ultrametric access patterns, and workload-aware energy scaling.

Cloud providers: JPCUB is the primary differentiator for compute instances. Providers compete on JPCUB per workload class. Customers select cloud providers based on JPCUB scores for THEIR specific workload mix.

Enterprise IT: Procurement is JPCUB-first. A server with higher TCO but significantly better JPCUB may win because total lifetime energy cost dominates hardware cost. TCO models are restructured: "energy per compute unit" is a primary line item, not a footnote.

Regulation: Mandatory JPCUB reporting for all data centers above a threshold size. Non-compliance penalties. JPCUB is a standard metric in corporate sustainability (ESG) reporting. Green bonds and sustainability-linked loans reference JPCUB improvement targets as performance indicators.

3. Competitive Landscape: Who Will Adopt First?

3.1 Cloud Providers

Provider	Likely Adoption Timeline	Rationale
Google Cloud	2028-2030	Strongest existing commitment to energy efficiency. Already publishes carbon-aware computing tools. TPU architecture provides unique data for JPCUB benchmarking.
AWS	2029-2031	Largest cloud provider by market share. Adopts metrics when customers demand them. Will implement JPCUB as an additional instance metadata field.
Microsoft Azure	2029-2031	Strong sustainability commitments. Carbon-aware SDK already exists. JPCUB complements existing Azure sustainability tools.

3.2 Hardware Vendors

Vendor	Likely Adoption Timeline	Rationale
AMD	2030-2032	Historically favorable efficiency metrics (Zen architecture). JPCUB advantages could differentiate from Intel.
Intel	2031-2033	Largest x86 vendor. Conservative metric adoption. Will adopt when customers demand it.

Vendor	Likely Adoption Timeline	Rationale
NVIDIA	2029-2031	GPU efficiency is a competitive differentiator for AI/ML workloads. JPCUB could demonstrate GPU advantage over CPU for specific workload classes.
ARM	2030-2032	big.LITTLE architecture is inherently efficiency-optimized. JPCUB could formalize what ARM already claims.

3.3 Regulatory Bodies

Body	Likely Adoption Timeline	Rationale
EU Commission (Energy Efficiency Directive)	2030-2032	Strongest regulatory driver. EU has historically led on energy regulation.
U.S. Department of Energy	2032-2035	DOE funds HPC; efficiency metrics are relevant for exascale computing programs.
China (MIIT)	2032-2035	Large data center market. Will adopt international standards if they align with domestic industry.

4. Technology Roadmap: From JPCUB to p-Adic Computing

4.1 The JPCUB → p-Adic Pipeline

JPCUB measures efficiency. But the METRIC implies the ARCHITECTURE. If JPCUB is the dominant competitive axis, hardware will be designed to optimize for it — and the mathematical structure that naturally captures hierarchical energy landscapes is p-adic.

Era	Metric	Architecture Implication
Era 1 (2025-2028)	JPCUB measures efficiency of existing hardware	No architectural change — measurement only
Era 2 (2028-2032)	JPCUB drives procurement decisions	Hardware vendors optimize existing architectures for JPCUB — better DVFS, workload-aware power management
Era 3 (2032-2035)	JPCUB is the dominant competitive axis	Novel architectures: p-adic number representations, ultrametric memory hierarchies, valuation-based energy scaling

4.2 The p-Adic Processor (Era 3)

A JPCUB-optimized processor in 2035 would have:

- Native p-adic arithmetic in ALU** — not IEEE 754 floating-point for error-bounded computation. p-adic numbers represent precision as valuation v_p , not mantissa bits. A computation that needs N digits of precision in $\mathbb{R}$ needs p^{-v_p} precision in $\mathbb{Q}_p$ — and v_p maps directly onto energy-per-operation.
- Hierarchical memory mapping** — memory access patterns that are naturally ultrametric. Cache hierarchies already approximate this (L1 → L2 → L3 → RAM → disk), but a p-adic processor makes this explicit: memory is addressed by p-adic valuation, and access latency is proportional to valuation gap.
- Valuation-based power gating** — compute units are powered on/off based on valuation requirements. A low-valuation (coarse) computation uses fewer gates at lower energy. A high-valuation (precise) computation

engages more gates. The energy budget is allocated by valuation, not by clock speed.

This processor would achieve $JPCUB = 10\times$ current best at iso-workload $\dagger$ — a design target derived from p-adic structural constraints, not incremental engineering improvements.

5. Risks and Caveats

1. **Cross-Architecture Validation Risk:** JPCUB's cross-architecture normalization is not yet validated (jpcub-validation project, 2026). If normalization fails for GPU/TPU/FPGA, JPCUB's scope contracts to CPU-only — less commercially impactful.
 1. **Metric Proliferation Risk:** If multiple efficiency metrics (JPCUB, Green500's next-generation metric, cloud-provider proprietary metrics) compete, none achieves standard status. The "metric that becomes the standard" scenario assumes convergence on one metric.
 1. **Regulatory Timing Risk:** Regulation moves slowly. If the EU Energy Efficiency Directive amendment is delayed past 2035, adoption in Eras 2-3 extends by 3-5 years.
 1. **Hardware Vendor Resistance:** If Intel or NVIDIA develops a proprietary efficiency metric and promotes it through marketing, JPCUB faces a standard-setting battle. The winner is not necessarily the "best" metric — it's the metric with the strongest institutional backing.
 1. **p-Adic Hardware Risk:** p-adic number representations in commodity processors require a manufacturing partner (TSMC, Samsung, Intel Foundry). If no partner commits by 2032, Era 3's p-adic processor remains a design concept.
-

6. Recommendations for QWAV

2026-2027: Metric Validation + Positioning

1. Complete JPCUB cross-architecture validation and publish as an open-access paper + dataset
2. Engage with cloud providers (Google, AWS, Azure) through their sustainability teams
3. Submit JPCUB as a proposed metric to the Green500 steering committee
4. Release open-source benchmark suite with reproducible methodology

2028-2029: Enterprise Adoption Catalysis

1. Publish JPCUB scores for all major cloud instance types (quarterly updates)
2. Partner with 1-2 Fortune 500 IT departments for procurement pilots
3. Submit JPCUB as a methodology to ISO/IEC JTC 1 (IT sustainability standards)
4. Publish annual "State of Computing Efficiency" report with JPCUB trends

2030-2035: Standardization + Architecture

1. Support regulatory adoption through technical working groups
 2. License JPCUB certification to hardware vendors (revenue model)
 3. Develop reference implementation of p-adic arithmetic in FPGA (technology demonstrator)
 4. Partner with at least one chip vendor for p-adic ALU prototype
-

7. The Alternative: What If JPCUB Fails?

If JPCUB fails to achieve standard status — because another metric wins, or because efficiency never becomes a dominant competitive axis — QWAV has two pivot paths:

Path A: p-Adic Computing Primitives. Abandon JPCUB as a metric and pivot to developing p-adic arithmetic libraries (software) and reference implementations (FPGA/ASIC) for high-performance computing. The value proposition shifts from "measure efficiency" to "provide efficient computation."

Path B: Consulting and Certification. Pivot from "create the metric" to "audit against the metric" — whichever metric wins, enterprises need independent verification of their efficiency claims. QWAV becomes the "Deloitte of computing efficiency" — auditing cloud-provider and hardware-vendor JPCUB (or equivalent) claims.

Both paths preserve QWAV's core asset: deep expertise in computing efficiency and p-adic mathematical methods.

Declarations

Funding: QWAV internal development. No external funding.

Conflicts of Interest: The author is the founder/principal of QWAV, which develops JPCUB and p-adic computing technology.

Forward-Looking Statements: This entire document constitutes forward-looking statements. See the disclaimer at the top of this document. Specific aspirational claims marked with † are design targets, not yet demonstrated.

Disclaimer: This document is for informational purposes only and does not constitute investment advice, a solicitation, or an offer to sell securities. Past performance and forecasts are not guarantees of future results.

Version History

Version	Date	Changes
v0.1	2026-07-31	Initial draft: 3-era forecast, competitive landscape, p-adic processor concept, pivot paths
