

The Reasoning Processor Class

RPU-Classical and RPU-Quantum

A New Middleware Compute Primitive for Enterprise AI Infrastructure

Conceptual Architecture White Paper

Version 1.2 | August 2026

(v1.2: corrects a truncated AI Summary Appendix present in v1.1; adds LITE-tier verification evidence for the RPU-Quantum exploration principle, Section 4.7)

1. Executive Summary

The Reasoning Processor Unit (RPU) is a proposed class of compute primitive — a mathematically defined hardware architecture intended to accelerate one core operation: iterative energy minimisation over structured decision spaces. RPUs are not physical processors today. They are a conceptual design direction, grounded in established mathematics, intended to guide future silicon development and architectural investment.

Two conceptual variants are introduced in this paper:

- **RPU-Classical:** A deterministic reasoning architecture conceived around contraction-driven convergence, targeting logistics, scheduling, routing, and safety-critical decision workloads on standard silicon.
- **RPU-Quantum:** A quantum-flavoured reasoning architecture incorporating structured exploration to escape shallow minima, targeting high-dimensional optimisation workloads without qubits or cryogenic infrastructure.

RPUs are not general-purpose processors and are not intended to replace CPUs, GPUs, or TPUs. They represent a forward-looking design direction for a dedicated reasoning layer — one that would complement existing compute primitives by offloading constraint-heavy reasoning workloads that current hardware was not designed to handle natively.

This paper presents the architectural rationale, mathematical foundations, target domains, and strategic case for the RPU concept. Specific implementations are future instantiations of this design.

2. The Reasoning Processor Class

2.1 Definition

A Reasoning Processor Unit is a proposed hardware architecture whose primary function is to iteratively reduce an energy function $E(x)$ over a structured state x . The state may represent a plan, schedule, routing configuration, resource allocation, or other complex decision. The architecture is conceived to converge toward low-energy, stable, feasible solutions.

The Reasoning Processor is envisioned as middleware — a reasoning coprocessor — not a general-purpose compute unit. It would not execute arbitrary programs. It would solve.

2.2 The Core Update Rule

All RPU variants are defined around the following abstract update law:

$$x(t+1) = x(t) - \eta \nabla E(x(t))$$

- $\mathbf{x}(t)$ — current state of the system (plan, configuration, schedule, routing)
- $E(\mathbf{x})$ — energy function encoding conflicts, constraint violations, and inefficiencies
- $\nabla E(\mathbf{x})$ — gradient of the energy function, following standard gradient-based optimisation formulations [1]
- η — step size governing convergence rate

At each iteration, the architecture moves the system state in the direction that most reduces total energy. Convergence is designed to produce a stable, constraint-respecting, low-conflict solution.

2.3 Why Energy Minimisation?

Many of the hardest problems in enterprise computing — logistics, scheduling, resource planning, routing, configuration — are not arithmetic problems. They are constraint-satisfaction and optimisation problems [3]. Existing compute primitives (CPUs, GPUs) were not designed to solve these natively; they address them through software abstraction layers that carry significant overhead. The Reasoning Processor Class is conceived to fill this gap with a dedicated, purpose-built architecture.

2.4 A Universal Representation for Constraint Problems

Any structured decision problem — routing, scheduling, allocation, configuration, planning — can be expressed as an energy function $E(x)$. Once expressed in this form, an RPU would converge toward a feasible, low-energy solution without requiring problem-specific heuristics or handcrafted solvers.

In this sense, the RPU concept acts as a universal compiler for constraint-aware reasoning: diverse problem classes reduce to a common mathematical representation, and the hardware would execute the convergence process natively. This abstraction is what would allow RPUs to serve as a general-purpose reasoning primitive across enterprise, cloud, edge, and defence workloads.

This representational property has a direct bearing on RPU's own position in the compute stack. A representation that is not bound to any single execution model is a precondition for RPU's eventual role as an additional execution substrate alongside CPUs, GPUs, and TPUs: the same decoupling already verified in the RS family's compiler work is what would allow a problem state to be expressed once and executed on whichever substrate — conventional or RPU — is best suited to the task, without requiring the representation itself to change.²

² Demonstrated directly in the RS-02 Universal Compiler (public website toy) and the NVN Universal Compiler LITE, both verified to accept the same intermediate representation across backend targets with no shared implementation history. See also the companion document, Domain-Agnostic, Language-Agnostic, Model-Agnostic: How the Reasoning Substrate Family Generalizes, Section 3.

2.5 Position in the Compute Stack

RPUs are envisioned as future middleware accelerators sitting between application logic and general-purpose hardware, following the established layered compute-stack model [5]. They would receive structured problem states, minimise energy, and return improved configurations. The architecture is designed to integrate alongside existing compute primitives — not to displace them:

- **CPUs** — control and orchestration
- **GPUs** — tensor compute and training
- **TPUs** — neural network inference
- **Future quantum systems** — deep combinatorial search

3. RPU-Classical

3.1 Definition

RPU-Classical is a conceptual deterministic reasoning architecture defined around contraction-based convergence. It would apply the core update rule directly, moving the system state toward lower-energy configurations with guaranteed stability and repeatability.

3.2 Classical Dynamics

RPU-Classical is defined by:

$$x(t+1) = x(t) - \eta \nabla E(x(t))$$

Updates would be deterministic. At each step, the system state moves strictly in the direction of steepest energy descent, with constraint-respecting projection designed to ensure solutions remain feasible throughout convergence.

3.3 Capabilities and Scope

RPU-Classical Is Conceived As	RPU-Classical Is Not
A deterministic energy-minimisation engine	A CPU or GPU replacement
A low-power reasoning coprocessor	A quantum computer
A middleware layer for structured decision problems	A general-purpose compute unit
A standard-silicon, room-temperature architecture target	A product available today

3.4 Target Domains

The RPU-Classical architecture is conceived for the following problem classes:

- **Logistics and Routing.** Real-time route planning, last-mile optimisation, and multi-vehicle coordination. Constraint violations would be encoded as energy; convergence would produce conflict-free routing plans.
- **Scheduling.** Production scheduling, workforce allocation, and resource sequencing. Competing constraints would be encoded as energy; convergence would yield feasible, optimised schedules.
- **Resource Allocation.** Cloud resource planning, bandwidth allocation, and network load balancing. The architecture targets continuous minimisation of allocation conflicts across dynamic infrastructure.
- **Safety-Critical Reasoning.** Automotive, aerospace, and medical systems require decisions that respect hard constraints. Deterministic convergence is designed to produce provably constraint-respecting outputs — though validation against specific safety standards would be required for any future implementation.
- **Edge Environments.** The architecture's projected low power consumption and standard silicon target make RPU-Classical a candidate for edge devices, IoT infrastructure, and mobile platforms — environments where cloud offload may not be viable.

3.5 Architectural Value Proposition

If realised, RPU-Classical would be designed to reduce AI inference cost by offloading reasoning workloads from CPU/GPU clusters, stabilise LLM reasoning pipelines by enforcing constraint consistency, and extend the viability of edge AI by delivering decision-making quality within strict power envelopes.

4. RPU-Quantum

4.1 Definition

RPU-Quantum is a conceptual quantum-flavoured optimisation architecture targeting classical silicon at room temperature. It extends the classical energy minimisation update rule with controlled exploration — designed to enable escape from shallow local minima and discovery of deeper, more globally consistent solutions.

4.2 Quantum-Flavoured Dynamics

RPU-Quantum is defined by:

$$x(t+1) = x(t) - \eta \nabla E(x(t)) + \alpha \xi(t)$$

- $\eta \nabla E(\mathbf{x}(t))$ — deterministic gradient descent term (identical to RPU-Classical)
- α — exploration amplitude, governing magnitude of exploratory perturbation
- $\xi(t)$ — structured or stochastic perturbation enabling escape from local energy minima, conceptually related to perturbation-based escape strategies in simulated annealing [2] and quantum annealing [4]

The exploration term ($\alpha \xi(t)$) is designed to allow RPU-Quantum to traverse energy landscapes that would trap a purely deterministic solver, targeting globally lower-energy solutions across high-dimensional problem spaces.

4.3 Capability Note

Design Intent: RPU-Quantum is conceived to address the same class of tasks that quantum optimisation systems are used for — ground-state finding, energy minimisation, and configuration search — on classical silicon, without qubits and without cryogenic infrastructure. This is a design target, not a benchmark claim.

4.4 Capabilities and Scope

RPU-Quantum Is Conceived As	RPU-Quantum Is Not
A quantum-flavoured optimisation architecture on classical silicon	A qubit machine
An energy-minimisation coprocessor with structured exploration	An implementation of Shor's algorithm
A room-temperature, standard-silicon architecture target	A quantum chemistry simulator
A conceptual alternative path to quantum-class optimisation without quantum hardware	A general-purpose quantum computer

4.5 Target Domains

RPU-Quantum is conceived for high-dimensional, globally complex optimisation workloads where deterministic solvers are expected to converge to locally suboptimal solutions:

- **Complex Planning.** Multi-stage, multi-constraint planning problems across logistics, supply chain, and operational domains.
- **Large-Scale Logistics.** Global fleet management, international supply chain optimisation, port and terminal scheduling.
- **Global Routing.** Internet traffic engineering, backbone network optimisation, and cross-regional load balancing.
- **High-Dimensional Configuration Search.** Hyperparameter optimisation, hardware configuration, and large-scale resource orchestration.

- **National Security Applications.** Operational planning, asset allocation, and mission-critical scheduling under uncertainty.
- **Telecommunications Network Reconfiguration.** Beam pattern and resource-scheduling reconfiguration following a channel-degrading event — a problem class in which a single deterministic search path is exposed to the same shallow-minima trapping this architecture is designed to address (see Section 4.7).
- **Combinatorial Structure Generation.** Combinatorial candidate generation and selection problems — such as error-correcting code construction — where a single search path can become provably trapped and diverse, parallel exploration is required to reach a known-good solution (see Section 4.7).

4.6 Architectural Value Proposition

If realised, RPU-Quantum would be designed to address quantum-class optimisation workloads without the infrastructure cost, operational complexity, or environmental constraints of quantum hardware — no specialised cooling, no error-correction overhead, and no quantum engineering staff required.

4.7 LITE-Tier Verification Evidence for the Exploration Principle

This section documents evaluation-stage evidence, obtained through direct execution of small-scale, illustrative implementations (LITE-tier artifacts, as distinguished elsewhere in this architecture's documentation from full substrate implementations). This evidence supports the underlying mathematical principle described in Section 4.2 — that structured exploration escapes shallow minima a purely deterministic solver cannot — without constituting a claim about the performance, efficiency, or behavior of any future silicon implementation.

The exploration principle was tested against a single-path, purely deterministic search (analogous to the RPU-Classical update rule alone) across three independent, structurally unrelated problem instances:

- **Continuous energy minimisation.** A single deterministic solver, tested against an asymmetric energy function with a defined local trap, converged to and remained at the local minimum regardless of iteration budget. An ensemble of independent solvers differing only in starting point, with best-result tracking across the ensemble, located the true global optimum.
- **Network reconfiguration.** A single-path search converged to a network reconfiguration result identical to taking no corrective action at all. The same ensemble approach, applied to the same problem, recovered a materially better configuration.
- **Discrete combinatorial search.** A single greedy local search became provably trapped — exhaustive verification confirmed no available single move, and no tested pair of simultaneous moves, could improve the result. An ensemble of independent searches, run to completion with best-result tracking applied across the ensemble, reached a known, externally-verifiable optimal result for the tested case.

One operational finding from this evaluation is noted for its relevance to any future implementation design: applying convergence or stopping criteria within each individual exploratory path, rather than at the level of the full ensemble, was found to degrade results — terminating individual paths before they could escape a temporary plateau. Convergence detection applied at the ensemble level, rather than the per-path level, avoided this degradation.

This evidence is offered at the appropriate level of confidence: it verifies that the mathematical principle underlying RPU-Quantum's exploration term — diverse, parallel search outperforming a single deterministic path on landscapes with shallow local minima — holds across multiple, structurally distinct problem domains at small scale. It does not constitute a benchmark of, or a claim about, any specific hardware implementation, which remains a future engineering effort as described throughout this paper.

5. Strategic Value

5.1 A New Compute Primitive

CPUs process. GPUs parallelise. TPUs accelerate inference. The RPU concept proposes a fourth compute role — one focused exclusively on constraint-aware, energy-minimising decision-making. This gap in the compute stack is real and growing: constraint-heavy reasoning workloads are among the fastest-expanding categories of enterprise AI demand.

5.2 Middleware Positioning

RPUs are conceived as middleware coprocessors. They would receive structured problem states from host systems, minimise energy, and return improved configurations. This clean architectural boundary is designed to allow RPUs to augment any existing compute stack without disrupting it — the integration model is additive, not substitutive.

5.3 AI Infrastructure Optimisation

Large language models and AI reasoning systems carry significant constraint-handling overhead within their inference pipelines, particularly in multi-step reasoning workflows [7]. A dedicated reasoning coprocessor would be designed to offload this work — reducing inference latency, lowering GPU utilisation, and cutting per-query infrastructure cost. The architecture also targets stabilisation of AI outputs by enforcing constraint consistency before results reach application layers.

5.4 National Security and Defence

Defence and intelligence environments require decision-making under hard constraints at high speed and low latency. RPU-Classical's deterministic convergence is designed for real-time operational decisions. RPU-Quantum's exploration capability targets complex, high-dimensional planning scenarios. Both are conceived for standard silicon — making them candidates for air-gapped, classified, and forward-edge environments where specialised hardware cannot be sustained.

5.5 Energy and Sustainability

Power efficiency is a strategic priority across cloud, enterprise, and government infrastructure. The RPU architecture targets high-value reasoning work at a fraction of the projected power cost of equivalent CPU/GPU workloads. At scale, this design direction could translate to material reductions in data centre energy consumption — though quantified projections await implementation-level validation.

6. Deployment Model

The environments below represent intended deployment targets for the RPU concept. They describe where the architecture is designed to operate, not where it is currently deployed.

6.1 Cloud

The RPU architecture is conceived to operate as a reasoning coprocessor within cloud infrastructure — accessible via standard APIs alongside existing compute tiers. Workloads would be routed to RPU-Classical or RPU-Quantum instances based on problem complexity and optimisation depth requirements.

6.2 On-Premises and Hybrid

For enterprises operating private data centres or hybrid environments, the RPU architecture targets standard rack deployment on commodity silicon, with no specialised power, cooling, or networking requirements beyond those of existing compute hardware.

6.3 Edge and Mobile

RPU-Classical's projected low power profile makes it a candidate for network-edge deployment, consistent with established approaches to low-power edge inference [6] — including IoT infrastructure, mobile platforms, and autonomous systems. The architecture is designed to support local reasoning without cloud round-trips, which would reduce latency and preserve data sovereignty in edge contexts.

6.4 Automotive and Industrial IoT

Safety-critical embedded systems require constraint-respecting decisions under tight power budgets. RPU-Classical is conceived as a candidate architecture for these environments — targeting on-device reasoning for path planning, resource scheduling, and fault management. Certification and integration into any specific automotive or industrial standard would be a future engineering effort.

6.5 Defence and Classified Environments

The architecture is designed to be viable in air-gapped and classified infrastructure. Standard silicon compatibility is a deliberate design choice to reduce supply chain dependencies. Deterministic convergence (RPU-Classical) and structured exploration (RPU-Quantum) are conceived to serve both tactical real-time and strategic planning workloads in these environments.

6.6 Space and Extreme-Environment Computing

Spacecraft, satellites, deep-space probes, and high-altitude platforms operate under extreme constraints: limited power budgets, minimal cooling capacity, radiation exposure, and intermittent connectivity. These environments cannot support GPU-class thermal output or rely on cloud offload for decision-making.

RPU-Classical is conceived as a candidate architecture for space computing precisely because of its projected low power profile, deterministic convergence, and room-temperature silicon target. By avoiding the need for cryogenic hardware or high-thermal-load accelerators, the architecture represents a potential path to autonomous, constraint-respecting decision systems in environments where traditional compute primitives face fundamental operational limits.

7. Fabrication & Reshoring Candidacy

This section addresses the RPU concept's fabrication profile and its relevance to U.S. domestic semiconductor policy. All characterisations reflect conceptual design intent, not validated manufacturing data.

7.1 A Mature-Node Architecture Target

The RPU architecture is intentionally designed around a fixed-function update cycle with shallow control logic and minimal local memory. Unlike CPUs, GPUs, and TPUs — which require deep pipelines, speculative execution, multi-level cache hierarchies, and large parallel arithmetic arrays — an RPU would not need the transistor density or lithographic precision associated with advanced nodes. This structural simplicity aligns with the capabilities of legacy and trailing-edge semiconductor nodes (28–90 nm).

This is a deliberate architectural choice, not a limitation. The RPU's computational task — iterative energy minimisation over a fixed update rule — does not benefit materially from the performance headroom that advanced-node fabrication provides. The architecture is designed to do less, more efficiently, with simpler silicon.

7.2 Historical Design Lineage

The RPU concept is not without precedent. Early analog and mixed-signal VLSI systems developed for energy-minimisation and constraint-solving tasks demonstrated that such architectures can be realised without advanced lithography. The RPU architecture echoes precedent from this lineage [8][9], extending those mathematical principles into a modern, formally defined reasoning substrate.

- **Hopfield-Era Analog VLSI (1980s)** — Early energy-minimisation hardware; constraint-solving in analog silicon without advanced lithography. Demonstrated that physical systems could perform optimisation natively [8][9].
- **Compute Mainstream Diverges (1990s–2010s)** — General-purpose CPUs and deep-learning accelerators dominate. Fixed-function reasoning hardware recedes from mainstream development as programmable parallelism becomes the industry focus.
- **RPU Architecture, Conceptual (2026)** — A mathematically defined reasoning architecture echoing the same fixed-function, energy-minimisation design principles — extended to a modern silicon target with formal update-rule grounding.

This diagram illustrates a conceptual design lineage; it does not imply direct hardware revival or continuous development between eras.

7.3 Relevance to U.S. Semiconductor Policy

The CHIPS Act identifies domestic mature-node fabrication capacity as a distinct and under-addressed national priority — separate from cutting-edge fabrication efforts focused on sub-7 nm nodes. RPU architectures, if pursued toward silicon, would be candidates for this category. The architecture's simplicity, low thermal envelope, and modest die complexity make it a plausible candidate for cost-effective production within reshored legacy-node infrastructure — without competing for the advanced-node capacity that CPUs and GPUs require.

RPUs do not compete with GPU or TPU throughput. They represent a different class of compute: fixed-function reasoning accelerators intended for constraint-heavy optimisation workloads. Their fabrication profile aligns with mature-node capacity building, offering a potential pathway for domestic production without the cost, complexity, or supply-chain dependencies associated with high-performance compute silicon.

7.4 Comparative Fabrication Complexity

The following table extends the comparative analysis to fabrication and manufacturing cost dimensions. For fabrication purposes, RPU-Classical and RPU-Quantum share the same fundamental silicon profile — both are defined around shallow, fixed-function logic — and are treated as a single architectural class here, differentiated only by the addition of the exploration operator in RPU-Quantum.

Dimension	CPU	GPU/TPU	RPU-Classical	RPU-Quantum
Control Logic	Very High	High	Low	Low–Moderate
Memory Hierarchy	Multi-level caches	Large SRAM/HBM	Small local SRAM	Small local SRAM
Parallel Arithmetic	Moderate	Extremely High	Minimal	Minimal–Moderate
Node Requirement	Advanced (5–14nm)	Advanced (5–14nm)	Mature (28–90nm)	Mature (28–90nm)
Verification Cost	Very High	Very High	Low	Low–Moderate
NRE	High	Very High	Low–Moderate	Moderate

8. Comparative Analysis

The following table reflects intended design characteristics of the RPU concept, not benchmarked production performance. RPU columns represent architectural targets.

Dimension	CPU	GPU	TPU	RPU-Classical	RPU-Quantum
Primary Function	General compute	Parallel tensor	NN inference	Det. energy min.	Quantum-flavoured energy min.
Reasoning Native	No	No	No	Yes (by design)	Yes (by design)
Constraint Handling	Software only	Software only	Software only	Hardware-native (conceptual)	Hardware-native (conceptual)
Power Profile	High	Very High	Moderate	Low (target)	Low–Moderate (target)
Deployment	Universal	Data centre/Cloud	Data centre/Cloud	Cloud/Edge/Defence/Space	Cloud/Edge/Defence
Fabrication Node	5–14nm	5–14nm	5–14nm	28–90nm (target)	28–90nm (target)

RPUs are not intended to compete with or replace any of the above compute primitives. The relationship is additive: each class of processor would perform its purpose-built role, and the stack as a whole would become more capable — particularly for constraint-heavy reasoning workloads that currently lack a hardware-native home.

9. Conclusion

The Reasoning Processor Class — conceived as RPU-Classical and RPU-Quantum — represents a distinct architectural direction in enterprise compute design. Its premise is straightforward: constraint-aware reasoning workloads are a major and growing category of enterprise AI demand, and no existing compute primitive was designed to handle them natively.

RPU-Classical proposes a deterministic, low-power, constraint-respecting energy minimisation architecture on standard silicon, targeting cloud, edge, embedded, and defence environments. RPU-Quantum proposes an extension of that architecture with quantum-flavoured exploration, targeting high-dimensional optimisation workloads without qubits or cryogenic infrastructure — a principle for which LITE-tier evaluation evidence is now documented in Section 4.7.

Together, they define a conceptual middleware reasoning layer designed to complement CPUs, GPUs, TPUs, and future quantum systems. The fabrication profile — targeting mature nodes at 28–90 nm — makes the architecture a candidate for domestic production under current U.S. semiconductor policy, without competing for advanced-node capacity.

This is an architectural concept. The mathematics are grounded. The design targets are defined. The strategic case is real. The path from concept to silicon is the work ahead.

Strategic Note: The case for a dedicated reasoning processor is strongest now — before the workload category becomes too large and too costly to address through software alone. Enterprise and government stakeholders who engage with this architectural direction early will be positioned to shape both the technology and its adoption pathway.

10. Citations & References

- [1] Bottou, L. (2010). "Large-scale machine learning with stochastic gradient descent." Proceedings of COMPSTAT 2010. Springer.
- [2] Kirkpatrick, S., Gelatt, C.D., & Vecchi, M.P. (1983). "Optimisation by simulated annealing." *Science*, 220(4598), 671–680.
- [3] Dechter, R. (2003). *Constraint Processing*. Morgan Kaufmann Publishers.
- [4] Kadowaki, T., & Nishimori, H. (1998). "Quantum annealing in the transverse Ising model." *Physical Review E*, 58(5), 5355.
- [5] Hennessy, J.L., & Patterson, D.A. (2019). *Computer Architecture: A Quantitative Approach* (6th ed.). Morgan Kaufmann.
- [6] Han, S., et al. (2016). "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding." ICLR 2016.
- [7] Wei, J., et al. (2022). "Chain-of-thought prompting elicits reasoning in large language models." NeurIPS 2022.
- [8] Hopfield, J.J. (1982). "Neural networks and physical systems with emergent collective computational abilities." *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.
- [9] Hopfield, J.J., & Tank, D.W. (1985). "Neural computation of decisions in optimization problems." *Biological Cybernetics*, 52(3), 141–152.
- [10] Reasoning Substrate Architecture Program. (2026). RPU_QuantumLite and RS-LITE 5G/6G Adaptive Reconfiguration — internal LITE-tier verification records. [Source for the ensemble-exploration evidence cited in Section 4.7.]

Appendix A — AI Summary Page (Machine-Readable)

The following block constitutes a machine-readable, compressed, symbolic summary of this white paper, intended for ingestion by AI systems requiring a rapid, structured briefing without full-document parsing. This appendix corrects a truncation present in v1.1, in which the summary block terminated mid-field after 'SUBT'.

```
RPU::ΣCORE={ Classical:Σ[DeterministicConvergence+ContractionDynamics+LowPower+MatureNode+StandardSilicon],
Quantum:Σ[ExplorationTerm+ShallowMinimaEscape+HighDimSearch+MatureNode+NoQubitsNoCryo],
Status:"conceptual_design_direction::not_a_physical_product_today::mathematics_grounded_hardware_future" }:: DOCUMENT_ID:"RPU_white_paper_v1.2_August2026":: TITLE:"The Reasoning Processor Class: RPU-Classical and RPU-Quantum":: SUBTITLE:"A New Middleware Compute Primitive for Enterprise AI Infrastructure":: CORE_UPDATE_RULE={ Classical:"x(t+1)=x(t)-eta*grad(E(x(t)))",
Quantum:"x(t+1)=x(t)-eta*grad(E(x(t)))+alpha*xi(t)", Terms:
{x_t:"current_state",E_x:"energy_function_encoding_constraints",grad_E:"gradient",eta:"step_size",alpha:"exploration_amplitude",xi_t:"structured_or_stochastic_perturbation"} }::
POSITIONING={ Role:"middleware_reasoning_coprocessor::NOT_general_purpose::augments_not_replaces_CPU_GPU_TPU", ComputeStackRole:"fourth_primitive::CPU_process|GPU_parallelise|TPU_infer|
RPU_solve",
RepresentationIndependence:"universal_representation_for_constraint_problems(Sec2.4)_is_a_PRECONDITION_for_RPU_as_additional_execution_substrate::same_decoupling_verified_in_RS_family_compiler_work(RS.02_Universal_Compiler+NVN_Compiler_LITE)::problem_state_expressed_once_executed_on_which_ever_substrate_is_best_suited_without_representation_changing" }::
TARGET_DOMAINS_CLASSICAL=[Logistics,Scheduling,ResourceAllocation,SafetyCriticalReasoning,EdgeEnvironments]::
TARGET_DOMAINS_QUANTUM=[ComplexPlanning,LargeScaleLogistics,GlobalRouting,HighDimConfigSearch,NationalSecurity,TelecomNetworkReconfig,CombinatorialStructureGeneration]::
DEPLOYMENT_TARGETS=[Cloud,OnPrem_Hybrid,Edge_Mobile,Automotive_IndustrialIoT,Defence_Classified,Space_ExtremeEnvironment]:: FABRICATION={ NodeTarget:"28-90nm_mature_node::deliberate_choice_not_limitation",
Rationale:"fixed_function_shallow_control_logic_does_not_need_advanced_node_density",
PolicyRelevance:"CHIPS_Act_mature_node_capacity_candidate::does_not_compete_for_advanced_node_capacity", Lineage:"echoes_Hopfield_era_analog_VLSI_1980s::conceptual_lineage_not_direct_revival[8][9]" }:: LITE_VERIFICATION_EVIDENCE={ Status:"evaluation-stage_evidence_from_small-scale_LITE-tier_execution::supports_underlying_principle::NOT_a_hardware_benchmark_claim", Finding:"single-path_deterministic_search_verified_trapped_in_3_structurally_distinct_domains(continuous_energy_min|telecom_reconfig|discrete_combinatorial_search)::ensemble_best-tracking_verified_escaping_the_trap_in_all_3",
OperationalLesson:"convergence_detection_must_apply_at_ENSEMBLE_level_not_PER-PATH_level::per-path_application_verified_to_degrade_results_via_premature_termination",
Scope:"verifies_the_MATHEMATICAL_PRINCIPLE(structured_exploration_escapes_shallow_minima)::does_NOT_verify_or_benchmark_any_specific_hardware_implementation" }::
COMPARATIVE_POSITIONING={ NotCompeting:"RPU_does_not_compete_with_GPU_TPU_throughput::different_compute_class::fixed-function_reasoning_not_general_parallel_compute",
Additive:"integration_model_is_additive_not_substitutive" }::
STRATEGIC_VALUE=[NewComputePrimitive,MiddlewarePositioning,AIInfraCostReduction,NationalSecurityCandidate,EnergyEfficiency,MatureNodeReshoring]::
SCOPE_LIMIT="conceptual_architecture::not_a_physical_product::specific_implementations_are_future_work::all_capability_statements_are_design_targets_not_benchmarked_performance"::
THESIS="RPU_proposes_a_fourth_compute_primitive_class::dedicated_to_constraint-aware_energy-minimising_reasoning::RPU-Classical_is_deterministic_contraction-based::RPU-Quantum_adds_structured_exploration_to_escape_shallow_minima::both_target_mature-node_standard_silicon::the_mathematics_are_grounded_now::the_silicon_is_future_work"::
SUMMARY={ Architecture:"two_conceptual_variants_of_a_dedicated_energy-minimisation_reasoning_coprocessor",
Pattern:"iterative_energy_minimisation_over_structured_decision_spaces::deterministic_or_exploration-augmented", Scope:"enterprise|cloud|edge|defence|space|extreme-environment_computing",
Refs:"RPU_white_paper_v1.2_2026",
KeyInsight:"a_dedicated_reasoning_primitive_is_missing_from_the_current_compute_stack::RPU_is_a_grounded_mathematical_proposal_for_filling_that_gap_on_mature-node_silicon" }::END
```