




◆ FOUNDATIONAL PAPER

The Quantum Civilization

Intelligence, Computation and the Architecture of a New
Civilizational Era

Steven Alber
AUTHOR · FOUNDER, QUANTUM CIVILIZATION

VERSION 1.1 · 28 OCTOBER 2026
QUANTUMCIVILIZATION.ORG

PUBLICATION RECORD

IDENTIFIER	QCRP-2026-001
SERIES	Quantum Civilization Research Papers
TITLE	The Quantum Civilization: Intelligence, Computation and the Architecture of a New Civilizational Era
TYPE	Foundational paper
VERSION	1.1 · published 28 October 2026
AUTHOR	Steven Alber
PUBLISHER	Quantum Civilization
HISTORY	1.1 · 28 October 2026 · Author attribution corrected: Steven Alber (author), Quantum Civilization (publisher). 1.0 · 28 October 2026 · First edition.
CITE AS	Alber, S. (2026). The Quantum Civilization: Intelligence, Computation and the Architecture of a New Civilizational Era. Quantum Civilization Research Papers, QCRP-2026-001, version 1.1. Quantum Civilization. https://quantumcivilization.org/publications/qcrp-2026-001/
ONLINE	https://quantumcivilization.org/publications/qcrp-2026-001/

Abstract

This founding paper sets out the hypothesis that the artificial intelligence revolution of our time may be remembered as the technological childhood of a larger transformation: the emergence of a quantum-enabled computational civilization. It separates demonstrated results from roadmaps and open questions across error correction, beyond-classical experiments, AI for quantum computing, autonomous science, the economics of quantum access and national strategy, as of October 2026. It argues that the first civilizational consequence of quantum computation, the global migration to post-quantum cryptography, is caused by a machine that does not yet exist, and that institutions now act on anticipated capability. The paper introduces four frameworks. The Civilizational Compute Stack is a five-layer model linking physical qubits to institutions. The constraint-migration hypothesis predicts that the binding constraint on quantum impact will move from physics to governance. Mandate Engineering is a discipline for the delegated, budgeted and verifiable use of scarce computation by autonomous agents. The Quantum Civilization Transition Model (QCTM) is a causal register of claims and indicators that replaces composite scores with refutable transitions. The paper assesses eight research frontiers and defines five pre-registrable experiments, led by a multi-provider study of shared agent budgets. It sets out the institution’s research standards and describes three conditional scenarios for 2027–2042, each with signposts and refutation signals.

KEYWORDS quantum computing · fault tolerance · autonomous agents · agentic AI · compute governance · Mandate Engineering · post-quantum cryptography · computational sovereignty · autonomous science · scenarios 2027–2042

The Quantum Civilization

	The Childhood of a Larger Transformation	4
PART I	• THE THESIS	5
01	The Central Hypothesis	6
02	Why Now: Three Curves Converging	8
PART II	• THE STATE OF THE FRONTIER	11
03	What Has Actually Been Demonstrated	12
04	Intelligence and the Quantum Machine	14
05	Autonomous Science and the Invoice of Discovery	14
06	The Economics of Scarce Computation	15
07	The Geography of the Quantum Age	16
PART III	• THE ARCHITECTURE OF A QUANTUM CIVILIZATION	18
08	The Civilizational Compute Stack	19
09	Mandate Engineering	20
10	Ownership and the Distribution of Value	22
11	QCTM: The Quantum Civilization Transition Model	23
PART IV	• THE RESEARCH PROGRAMME	26
12	Eight Research Frontiers	27
13	The First Experimental Programme	28
14	The Standard We Hold Ourselves To	29
PART V	• HORIZONS	31
15	Beyond the Classical World: Three Futures, 2027–2042	32
16	The Questions That Define the Territory	33
	Earning the Ambition	35
	Method and Evidence Standard	35
	Sources	36

The Childhood of a Larger Transformation

Every technological era is remembered for its machines and decided by its institutions. The steam engine is remembered; the factory system, the joint-stock company and the labour law it forced into existence did the deciding. The microprocessor is remembered; the platform, the data centre and the terms of service decided who would hold power in the age it opened.


We are living through the most visible technological transition of the century. Artificial intelligence has moved, within a few years, from research laboratories into the daily operation of science, commerce and government. It is natural to assume that this is the main event.

Quantum Civilization begins from a different possibility. The artificial intelligence revolution of our time may eventually be remembered as the technological childhood of a much larger transformation: the emergence of a **quantum-enabled computational civilization**. In such a civilization, an increasing share of discovery, coordination and decision is carried out by autonomous systems, and those systems draw on a new kind of computation. It is not merely faster than the computation that came before. It is different in kind, capable of representing nature at the level where nature itself is quantum mechanical.

This is a research hypothesis, not a prophecy. Its value lies in what it forces us to ask. Imagine autonomous intelligence, fault-tolerant quantum computation, robotic laboratories and synthetic economic systems working as one infrastructure. Who would authorize that infrastructure to act? Who would pay for its computation, verify its results and own its discoveries? How would the balance of power shift among states, firms, laboratories and citizens?

These questions are usually treated as distant and secondary, to be settled after the engineering. This paper argues the opposite. The engineering is already producing results that change the timetable. Error-corrected quantum memories now improve as they are made larger [1][51]. Governments have begun rebuilding the cryptography of the world economy against a quantum computer that has not yet been built [29][30]. Autonomous agents already plan experiments, call laboratory instruments and submit computational jobs through standard protocols [17][18][21]. The institutional architecture of the quantum age is not being designed later. It is being designed now, mostly by default and largely out of public view.

Quantum Civilization exists to make that design visible, measurable and contestable. This founding paper sets out what is established, what is plausible and what remains genuinely open. It introduces the conceptual frameworks the institution will develop and test. And it defines the first experiments through which an ambitious thesis can begin to earn its evidence.



PART I

The Thesis

CHAPTER 01

The Central Hypothesis

A civilization defined by its computation

Historians describe civilizations by their dominant capacities. Agricultural civilizations organized themselves around the capacity to store surplus. Industrial civilizations were organized around energy converted into motion. The civilization now taking shape is organizing itself around a newer capacity: computation that can perceive, decide and act.

We use the term **computational civilization** for a society in which three conditions hold at once:

1. A growing share of discovery, coordination and decision-making is performed or mediated by autonomous computational systems rather than by people directly.
2. Access to computation, including its hardware, energy, data and the legal right to deploy it, becomes a primary determinant of economic and institutional power.
3. The rules that govern computation, such as who may run it, on whose behalf, at what price and with what accountability, become constitutional questions in the practical sense. They shape the distribution of power as decisively as property law or the franchise once did.

By these criteria the transition is already under way on classical hardware. The question this institution investigates is what happens when that civilization acquires a second, fundamentally different engine.

Why quantum computation is different in kind

Classical computers manipulate bits that are either zero or one. Quantum computers manipulate quantum states, which can be placed in superposition, entangled with one another and made to interfere. Popular accounts often claim that a quantum computer “tries every answer at once”. This is wrong, and the error obscures what is genuinely powerful. A quantum algorithm works by arranging interference so that the amplitudes of wrong answers cancel and the amplitudes of useful answers reinforce. Only problems with the right mathematical structure allow this to be done efficiently.

The problems that do have that structure matter enormously. In 1982 Richard Feynman observed that nature at the smallest scales is quantum mechanical, and that simulating it on classical machines becomes exponentially expensive as systems grow. In 1994 Peter Shor showed that a quantum computer could factor large numbers efficiently, which would break the public-key cryptography that protects most of the world’s digital communication. Between these two landmarks lies the core promise of the field: the simulation of molecules, materials and quantum matter beyond classical reach, and a small but consequential set of algorithmic speed-ups.

For most tasks a quantum computer offers no advantage at all. Whether quantum machines will accelerate the training of large AI models is, on the present evidence, unproven [11][12]. That is precisely why the civilizational question is interesting. Quantum computation will not replace classical computation. It will join it as a scarce, specialized and expensive resource, consumed inside hybrid workflows for the problems where it matters. Increasingly those workflows will be orchestrated by autonomous systems.

The hypothesis in three parts

The central hypothesis of Quantum Civilization can be stated as three linked claims, each of which can be investigated separately:

H-I · Capability. Fault-tolerant quantum computation will make specific classes of scientific and economic problems tractable that are intractable for classical machines, beginning with the simulation of quantum systems.

H-II · Integration. These capabilities will be consumed mainly by autonomous systems working inside hybrid classical-quantum workflows, not by human users directly. The typical quantum job of the 2030s will be requested by an agent, priced by a cloud, executed by a device and validated by another machine.

H-III · Institution. The civilizational consequences of quantum computation will depend less on physics than on who is authorized to commission it, who pays for it, who verifies it and who owns what it produces.

The first claim belongs to physics and computer science. The second belongs to systems engineering and AI. The third is where this institution intends to make its distinctive contribution. It can be compressed into a single sentence that serves as the institution's working principle:

Physics will decide what becomes computable. Institutions will decide what becomes of it.

The three claims are deliberately separable. H-I could be confirmed while H-III proves irrelevant, for example if quantum advantage arrives but is absorbed by existing institutions without consequence. H-III could matter even if H-I is delayed for decades, because the institutional patterns of autonomous computation are being set now on classical machines. A research programme built on separable claims can make progress on each without betting everything on one.

The classical future that still happens

An honest thesis must state its strongest counter-thesis. Ours is this: the decisive transformation of the coming decades may be accomplished entirely by classical AI running on classical hardware, with quantum computation remaining a scientific instrument of limited economic reach.

This is not a fallback position to be mentioned and discarded. It is one of the scenarios this institution will track with the same seriousness as the others (Part V). It also sharpens the research agenda. If the institutional questions of H-III are real, they are real in both futures. Who authorizes an autonomous system to spend computational resources, and who owns what it discovers, are questions for a classical agentic economy as much as for a quantum one. Quantum computation intensifies them. Scarcity, queueing, calibration drift, shot-based pricing and the irreversibility of executed jobs make every weakness in the governance of delegated computation more costly. The research is therefore valuable on any timeline, and its quantum-specific claims can be tested rather than assumed.

CHAPTER 02

Why Now: Three Curves Converging

Ambitious hypotheses become research programmes when the evidence begins to move. In 2026 three independent curves are moving at once.




FIGURE 1 Three curves converging, 2023–2026. A conceptual diagram, not a plot of data.

The first curve: error correction crosses the threshold

For three decades the central obstacle to useful quantum computation has been noise. Physical qubits lose their quantum information quickly, and every operation adds error. Quantum error correction answers this by encoding a single *logical* qubit across many physical qubits. That only works if the physical error rate is below a threshold. Below it, adding more qubits suppresses logical errors. Above it, adding qubits only adds more ways to fail.

In 2024 Google Quantum AI reported a distance-7 surface-code quantum memory built from 101 physical qubits, with a logical error of 0.143% per correction cycle. Each step up in code distance cut the logical error rate by a factor of 2.14, and the team also demonstrated real-time decoding on a separate processor [1]. This is the scaling behaviour that fault-tolerant computing requires, now demonstrated in a real device. It is a result about quantum *memory*, and the path from protected memory to universal logical computation still contains substantial engineering. But the direction of travel is now an experimental fact rather than a theoretical promise.

The curve has kept rising. In October 2026 Quantinuum reported a logical memory fidelity of 99.9995% on its trapped-ion Helios system, roughly five errors in a million, using a colour code [51]. The same month, Google reported in *Science* that *magic-state cultivation* on its Willow processor reduced errors in prepared magic states about forty-fold [52]. Magic states are the costly ingredient that universal fault-tolerant computation requires. Neutral-atom systems built by Harvard, QuEra and collaborators have executed operations on reconfigurable logical qubits [5]. Photonic approaches have demonstrated manufacturable components and networked modules [6][7]. Silicon-based qubits have cleared manufacturing milestones [39]. And reinforcement-learning controllers now steer error correction as hardware drifts [10].

Not every platform claim has survived scrutiny. Microsoft's topological programme must be read alongside an independent 2026 analysis questioning how topological signatures were interpreted [8]. That is not a weakness of the field. It is the field working as science should.

The second curve: autonomous systems acquire hands, tools and budgets

The second curve belongs to artificial intelligence. In 2023 the Coscientist system combined a large language model with laboratory automation to plan and execute chemistry experiments [17]. Berkeley's A-Lab automated the synthesis of inorganic materials [18]. In January 2026 an author correction revised how some of its compounds had been identified and described as new [19], a reminder that verification is part of the infrastructure of autonomous science, not an afterthought. Multi-agent laboratory systems with built-in self-correction followed in 2026 [20].

These systems are no longer confined to chemistry benches. A 2026 IEEE paper described a Model Context Protocol server that lets AI agents submit and manage quantum jobs in hybrid quantum-HPC environments [21]. A protocol reported as Q-LEASE proposes authorization that follows a quantum job from queue admission through execution to the release of its result [22]. Agents are acquiring the ability to commission computation, including quantum computation, on someone's behalf. The governance of that ability is barely formed.

The third curve: institutions move before the machine arrives

The third curve is the most revealing. In August 2024 the U.S. National Institute of Standards and Technology approved its first three post-quantum cryptography standards, FIPS 203, 204 and 205 [29]. In October 2026 it finalized IR 8547, its guidance for the transition away from quantum-vulnerable public-key cryptography [30]. In May 2026 the OECD Council adopted a Recommendation on Quantum Technologies [46]. National programmes now span Europe, the United States, China, Japan, South Korea, India, Singapore, the United Kingdom, Australia and Canada [31][33][35][36][37][38][40][41].

The cryptographic migration deserves particular attention, because it reveals something new about how technology reshapes institutions:

INSIGHT


The first civilizational consequence of the quantum computer is being caused by a quantum computer that does not yet exist. Data encrypted today can be stored and decrypted later, so the mere *expectation* of capability forces institutions to reorganize in advance. Quantum computation is the first technology whose institutional effects are arriving on a schedule set by its anticipated future, not by its present performance.

The pattern will not be confined to cryptography. Expectations of quantum capability are already shaping national investment, export policy, talent flows and corporate strategy [31][32]. If institutions act on anticipated capability, then understanding how capability actually develops becomes a matter of public consequence. Distinguishing a demonstrated result from a company roadmap, and a benchmark from a breakthrough, is no longer a technical courtesy. It is a civic necessity.

Where the curves meet

Each curve is significant on its own. Their convergence is what defines the territory of Quantum Civilization. Error-corrected quantum computation is becoming an engineering programme. Autonomous systems are acquiring the tools and the authority to commission computation. And institutions are reorganizing around anticipated capability before the capability exists. The meeting point of the three is the place where quantum capability is commissioned, governed and turned into consequence, and no existing discipline owns it. Physicists study the devices. AI researchers study the agents.

Economists, lawyers and political scientists study institutions. The architecture that will connect them is being assembled without a dedicated science. This institution intends to supply one.



PART II

The State of the Frontier

CHAPTER 03

What Has Actually Been Demonstrated

A research institution that studies the consequences of a technology must first be exact about the technology. This chapter separates three kinds of statement that public debate routinely blurs: what has been **demonstrated** in peer-reviewed experiments, what has been **announced** in company roadmaps, and what remains an **open question**. The distinction is not a matter of caution. It is the raw material of analysis.

Error correction: from theory to scaling law

The most consequential result of the past three years is the experimental arrival of below-threshold error correction. Google's surface-code memory showed logical error rates falling as the code grew, with a suppression factor of 2.14 per step in code distance and a logical error of 0.143% per cycle at distance seven [1]. Quantinuum has since reported a logical memory fidelity of 99.9995% using a colour code on its trapped-ion Helios system [51]. Neutral-atom arrays have run circuits on reconfigurable logical qubits [5], and Quantinuum and Microsoft reported a joint logical-qubit demonstration in 2024 [43].

Two further results of 2026 show the engineering stack maturing around these memories. Reinforcement-learning controllers have been used to stabilize error correction under hardware drift [10]. Magic-state cultivation on Willow cut the error of prepared magic states about forty-fold [52], reducing one of the largest projected overheads in fault-tolerant architectures.

The precise boundary of these results matters, because it marks where the next decade of work lies. Below-threshold *memory* is demonstrated. Large-scale *universal logical computation*, meaning long sequences of error-corrected logical gates on many logical qubits, is the next frontier. Between the two lie the cost of logical operations, decoding speed, the production of magic states and the physical-qubit overhead per logical qubit.

Beyond-classical experiments and the moving classical frontier

Quantum hardware has repeatedly performed specific computational tasks that challenge the best classical methods. IBM's 127-qubit experiment in 2023 estimated observables of noisy circuits at a scale that strained brute-force classical simulation [4]. China's 105-qubit Zuchongzhi 3.0 set a new random-circuit-sampling benchmark in 2025 [33]. Google's 2025 *Quantum Echoes* experiment observed constructive interference at the edge of quantum ergodicity, a measurement with direct scientific meaning for the physics of quantum chaos [3].

Each of these results has the same structure. It is a challenge issued to classical computer science, and classical computer science answers. Improved tensor-network and simulation methods have repeatedly narrowed earlier claims of advantage, as the Chinese Academy of Sciences itself noted in presenting Zuchongzhi 3.0 [34]. Ewin Tang's 2018 "dequantization" of a celebrated quantum recommendation algorithm showed that a speed-up can disappear once the classical competitor is given the same access to data [12].

INSIGHT



Quantum advantage is not a finish line. It is a contest. Every claim is provisional against the strongest classical method that exists, and that method keeps improving. The meaningful question is never “has quantum advantage been achieved?” but “for which task, against which classical baseline, at what total cost, and for how long?”

This framing turns a source of public confusion into a research instrument. A rigorous account of quantum progress must track not only the quantum machines but the classical frontier they are measured against. Benchmarking must also be methodologically sound, which is itself an active research field [2].

Platforms: many roads, no settled winner

Superconducting circuits (Google, IBM), trapped ions (Quantinuum, IonQ, university groups including Waterloo’s IQC), neutral atoms (QuEra and Harvard–MIT collaborators), photonics (PsiQuantum, Xanadu), silicon spin qubits (UNSW and Diraq) and topological approaches (Microsoft) each pursue different trade-offs between fidelity, speed, connectivity and manufacturability [5][6][7][39][44][48]. Company roadmaps, such as IBM’s plan for a large fault-tolerant system by the end of the decade, state engineering intentions, not achieved results [49].

For a civilizational analysis, the plurality of platforms is itself significant. Competing physical architectures mean competing supply chains, national programmes and standards. The geography of quantum capability will be shaped by which platforms scale, and no platform has yet established a decisive lead.

The ledger of the frontier

TABLE 1 The frontier in October 2026, separated by evidence class.

Domain	Demonstrated	Announced or projected	Decisive open question
Error correction	Below-threshold logical memory; 99.9995% logical fidelity in memory [1][51]	Large-scale fault-tolerant systems by 2029–2030 [49]	Cost and speed of universal logical gates at scale
Beyond-classical tasks	Specific sampling and physics tasks challenging classical methods [3][4][33]	Broad commercial advantage	Which economically relevant task survives the best classical baseline
Logical architecture	Logical-qubit operations on several platforms [5][43]; magic-state cultivation [52]	Modular, networked machines [6][7]	Physical-qubit overhead per useful logical operation
AI for quantum	RL control of error correction [10]; AI across the quantum stack [9]	AI-designed codes and calibration at scale	How far AI can compress the engineering timeline
Quantum for AI	Theoretical speed-ups for specific structured problems	Quantum-accelerated machine learning	Any advantage that survives data-loading costs and classical competition [11][12]

CHAPTER 04

Intelligence and the Quantum Machine

Public imagination expects quantum computers to make artificial intelligence more powerful. The evidence points to a more interesting and less obvious relationship.

The first fusion runs from AI into quantum

The clearest demonstrated contribution so far runs in the opposite direction: **artificial intelligence is helping to build the quantum computer before the quantum computer can help artificial intelligence**. Machine learning is used across the quantum stack, from device calibration and pulse design to error decoding [9]. In 2026 reinforcement learning was used to steer error correction itself as hardware conditions drift [10]. Decoding, calibration and control are exactly the bottlenecks that stand between today's quantum memories and tomorrow's logical processors.

This reverses the usual causal story, and it has a strategic implication. The organizations with the strongest AI capabilities are positioned to compress the engineering timeline of quantum computing. Capability in one domain becomes leverage in the other. The coupling of the two technologies is already a fact of industrial organization, well before it is a fact of computational performance.

What quantum computation can and cannot do for AI

Quantum machine learning is an active and contested field. The decisive constraints are well understood. Classical data must be loaded into quantum states, which can erase theoretical gains. Classical models that learn from data can rival quantum models even where computing the target function classically is hard [11]. Speed-up claims must be tested against the strongest classical competitor [12]. No general quantum speed-up for training large language models has been demonstrated.

The plausible contribution of quantum computation to intelligence is therefore specific rather than general. It lies where the data are themselves quantum: molecular, material and physical systems whose behaviour classical models can only approximate. An autonomous research agent that can commission an accurate quantum simulation of a catalyst or a battery material is not "smarter" in any general sense. It is better informed about a part of nature that was previously beyond computational reach. That is the form in which quantum capability is most likely to enter the cognition of machines: as a new *sense organ* for the quantum world, attached to classical intelligence.

CHAPTER 05

Autonomous Science and the Invoice of Discovery

From automated instruments to autonomous inquiry

Science has used automation for decades. What is new is the delegation of *inquiry*: systems that formulate hypotheses, design experiments, operate instruments, interpret results and decide what to

do next. Coscientist planned and executed chemical syntheses using a language model connected to laboratory hardware [17]. Berkeley's A-Lab ran an automated pipeline for synthesizing inorganic materials [18]. Multi-agent laboratory systems with self-correction followed in 2026 [20]. Agents can now submit and manage jobs on hybrid quantum–HPC infrastructure through standard tool protocols [21].

The 2026 correction to the A-Lab paper, which revised how compounds were identified and described as new [19], is not a footnote to this story. It is one of its central lessons. When discovery is automated, verification must be automated with equal seriousness, or autonomous science will produce errors at the speed of machines and corrections at the speed of humans.

When scientific discovery has an invoice

Autonomous science turns discovery into a sequence of authorized, metered actions. Each hypothesis tested by an agent consumes compute time, instrument time, reagents and, increasingly, quantum processor time. Each of these has a price and an owner, and each requires an authorization. The scientific record of an autonomous discovery is therefore inseparable from its *transaction record*. Somebody permitted the agent to spend. Somebody paid. Somebody's infrastructure executed the work. Somebody will claim the result.

INSIGHT



In autonomous science, the provenance of a discovery and the accounting of its cost become the same chain of evidence. An institution that cannot reconstruct who authorized, funded and verified a machine-made discovery cannot credibly claim to know that the discovery is real, or whom it belongs to.

This is one of the founding observations of Quantum Civilization's research programme. It is why the institution treats authorization, budgeting and verification not as administrative matters but as part of the epistemology of the coming era.

CHAPTER 06

The Economics of Scarce Computation

Quantum computation already has a price list

Before quantum computation has a mass market, it already has a pricing architecture. Amazon Braket charges per task, per shot or per reserved device time, depending on the hardware, with classical cloud services billed separately [26]. Its identity and access system defines explicit permissions for creating and cancelling quantum tasks [27], and its service quotas bound what any account can do [47]. IBM documents a fair-share scheduler for allocating scarce processor time among users [25], and its public Open plan offers up to ten minutes of quantum processor time in a rolling 28-day window [28].

These details matter because units of account shape behaviour. When computation is priced by the shot, the task or the reserved second, the design of experiments, the competitiveness of research

groups and the strategies of autonomous agents all bend around those units. The quantum economy is already being given its grammar.

INSIGHT



Metered scarcity is the defining economic condition of early quantum computation.

Whoever defines the units, prices and queues of quantum access defines who can participate in the science and commerce built on it.

Advantage must be measured in total cost, not in device time

Economic claims for quantum computation are most often made for optimization and finance. The evidence demands precision. An independent 2025 study of quantum annealing found favourable results on some test problems, while in a power-generation scheduling case the hybrid solver did not outperform the classical Gurobi solver in either runtime or solution quality [13]. The Quantum Optimization Benchmarking Library, published in 2026, provides ten problem classes with open data and code. That gives the field a common yardstick, which is the precondition for any credible claim of advantage [14][15]. In finance, rigorous resource estimates for pricing options under stochastic volatility show what fault-tolerant hardware would be required, rather than an advantage available today [16].

The lesson is methodological and it is decisive. A hybrid quantum workflow contains classical preprocessing, compilation, queueing, execution, repetition, measurement and classical post-processing. Its value must be judged on the *whole workflow*: wall-clock time, total cost, solution quality and the strength of the classical alternative. The duration of a quantum operation alone says almost nothing.

CHAPTER 07

The Geography of the Quantum Age

Computational sovereignty

Quantum technology has become an instrument of national strategy. The OECD's 2025 overview documents the spread of national quantum strategies and policies [31]. UNESCO's 2026 global report examines quantum technology as a question of global development and access [32]. In May 2026 the OECD adopted a Council Recommendation on Quantum Technologies, bringing quantum governance into the formal instruments of international policy [46].

The landscape is broad and asymmetric:


- **Europe** has adopted a quantum strategy and is integrating quantum processors with its high-performance computing infrastructure [41].
- **The United States** combines federal standard-setting, notably NIST's cryptographic leadership [29][30], with the largest concentration of private quantum companies.
- **China** has demonstrated large superconducting processors, including the 105-qubit Zuchongzhi 3.0 [33][34].

- **Japan** has paired a 256-qubit superconducting system built by RIKEN and Fujitsu with the integration of quantum computation and the Fugaku supercomputer for quantum chemistry [35].
- **South Korea** is building quantum research capacity through KAIST and the national standards institute KRISS [36].
- **India's** National Quantum Mission spans hardware, software, communication and applications from 2023 to 2031 [38].
- **Singapore's** Centre for Quantum Technologies anchors a national ecosystem, and its National Quantum Office has announced a strategic partnership with Quantinuum [37].
- **The United Kingdom's** National Quantum Computing Centre operates testbeds across multiple qubit platforms [40].
- **Australia and Canada** host leading silicon and trapped-ion research [39][48].

National goals are not achieved dominance, and a strategy document is not a capability. The pattern nevertheless supports a concept this institution will develop: **computational sovereignty**. It means the capacity of a state, region or institution to access, govern and benefit from frontier computation on its own terms. In a computational civilization, sovereignty is measured not only in territory and currency but in access to processors, energy, data, talent and standards. Quantum computation adds a new and highly concentrated layer to that measure.



QUANTUM
CIVILIZATION
RESEARCH INSTITUTE



PART III

The Architecture of a Quantum Civilization

The Civilizational Compute Stack

Five layers between a qubit and a civilization

Debates about the impact of quantum computing usually jump from a hardware milestone straight to a civilizational consequence: a new processor is announced, and commentators speak of revolutions in medicine, finance and security. Everything in between is skipped. Those intermediate layers are where consequences are actually made, delayed or prevented.

Quantum Civilization proposes a five-layer model, the **Civilizational Compute Stack**, as the basic map of its research:




FIGURE 2 The Civilizational Compute Stack and the constraint-migration hypothesis.

TABLE 2 The Civilizational Compute Stack.

Layer	What it contains	Typical evidence
L1 · Physical capability	Qubits, fidelities, coherence, connectivity, calibration	Device benchmarks, error rates under stated conditions
L2 · Logical computation	Error-corrected qubits and gates, decoding, magic states	Logical error per cycle, logical gate fidelity, overhead
L3 · Hybrid workflow	Classical-quantum-AI pipelines, scheduling, pricing	End-to-end time, cost and quality against classical baselines
L4 · Autonomous agency	Agents that commission computation under mandates and budgets	Authorization records, budget integrity, verification chains
L5 · Institutions	Science, markets, law, states and the distribution of value	Adoption, contracts, ownership, policy, sovereignty

Each layer has its own physics, engineering, economics or politics, and its own evidence. A breakthrough at one layer becomes a consequence at the next only if a specific transition succeeds. Below-threshold memory (L1–L2) becomes economic value only if a workflow outperforms its classical alternative at acceptable cost (L3), if autonomous systems are authorized and able to commission it (L4), and if institutions adopt, finance and distribute its results (L5).

The constraint-migration hypothesis

The model yields a hypothesis that is, as far as this institution can determine, not stated elsewhere in this form:

The constraint-migration hypothesis. As quantum computation matures, the binding constraint on its civilizational impact moves *upward* through the stack: from physics (L1–L2), to workflow economics (L3), to the governance of autonomous agency (L4), to institutional design (L5). The scarce resource of the late quantum era will not be qubits. It will be trustworthy authority to deploy them.

The hypothesis has observable implications. If it is correct, research effort, investment and policy attention should shift over time from hardware toward orchestration, authorization and institutional design. The gap between what is technically possible and what is institutionally deployable should widen before it narrows. The first serious failures of the quantum era should then be failures of governance: misallocated computation, unverifiable results and disputed ownership, rather than failures of physics. Each implication can be tracked. The institution’s Observatory and QCTM framework are designed to track them.

CHAPTER 09

Mandate Engineering

Who authorizes the quantum job?

Consider an ordinary scenario of the near future. A pharmaceutical company authorizes an autonomous research agent to explore candidate catalysts. The agent spawns sub-agents. They purchase classical simulation time from one cloud, GPU time from another and quantum processor time from a third, each priced in different units. A calibration drift forces a resubmission. Two sub-agents unknowingly submit the same job. A quantum task completes after the company has revoked the agent’s authority. The result points toward a valuable discovery.

Who authorized each expenditure? Did any agent exceed its mandate? Can a completed quantum computation be “cancelled”? Who pays for the duplicated job? Who owns the discovery? Was the quantum component actually necessary, or would a classical method have done as well?

None of these questions is exotic. Each arises from the combination of autonomous agency, multi-provider computation and metered scarcity that the evidence in Part II documents. Existing systems answer pieces of the problem. Cloud identity systems define who may create and cancel tasks [27]. Schedulers allocate scarce processor time [23][24][25]. The reported Q-LEASE protocol binds authorization to the lifecycle of an individual quantum job [22]. What is missing is a coherent discipline

for the delegated, budgeted and verifiable use of scarce computation across agents, organizations and providers. Quantum Civilization calls that discipline **Mandate Engineering**.

Six concepts that must never be merged

Mandate Engineering begins with conceptual precision. Much confusion in debates about autonomous agents comes from collapsing distinct concepts into one:

1. **Principal.** The person or organization that grants a mandate, in the sense of the relevant legal relationship.
2. **Agent identity.** The means of distinguishing which software component made a request. Identity says *who acted*. It does not say on whose behalf the agent may act.
3. **Organizational authority.** The competence assigned by law and internal governance to decide on an organization's behalf.
4. **Mandate.** The binding of a specific actor to a permitted purpose, operations, resources, duration and constraints.
5. **Technical credential.** A token or key that implements a mandate. Possession of a credential does not by itself prove legitimate authority.
6. **Liability.** Responsibility for outcomes. It does not arise automatically in an agent. It rests on agreements among people and organizations within a legal context.

The minimal mandate record

From these distinctions follows a minimal record that any governable delegation of computation should carry:

TABLE 3 The minimal mandate record proposed by Mandate Engineering.

Field	Purpose
Principal and issuing authority	Who grants the mandate, and whether the issuer is competent to grant it
Agent	Which identified component may act
Purpose and permitted operations	What the computation is for and what it may do
Workflow and data class	Which pipeline, and what sensitivity of data
Permitted providers and regions	Where computation may run
Validity and expiry	When the mandate begins and ends
Delegation rule	Whether and how the mandate may be passed on
Budget	Ceilings in currency and in resource units (shots, tasks, processor-seconds)
Reservation rule	How funds are committed before execution
Stop conditions	When execution must halt
Result destination	Where outputs go and who may access them
Evidence	The record that allows independent reconstruction

Two principles govern the structure. **Monotonic attenuation:** every delegated mandate is no broader than its parent, in purpose, budget, duration and permissions. **Irreversibility of execution:** a computation that has run cannot be revoked out of existence. Revocation can stop future actions, but the record, the cost and the consequences of completed work must be preserved and reconciled.




FIGURE 3 A delegation tree under Mandate Engineering. Every child mandate is no broader than its parent, and one shared ledger enforces the budget across agents and providers.

Why quantum computation sharpens the problem

The problems of Mandate Engineering exist for any expensive computation. Quantum computation sharpens them in identifiable ways: long and variable queues; calibration that changes between submission and execution; billing by shots and tasks rather than time; execution that cannot be paused or rolled back; and workflows that span several providers with incompatible units of account. Whether these differences make quantum workflows *qualitatively* different, or only quantitatively harder, is itself one of the institution's first empirical questions. The experimental programme in Part IV tests the same mechanisms on classical and quantum workloads side by side.

CHAPTER 10

Ownership and the Distribution of Value

Who owns what autonomous systems discover?

When an autonomous agent commissions a quantum simulation that leads to a new material, the value of the discovery passes through many hands. There is the principal who funded the agent, the provider of the model, the operator of the quantum processor, the owner of the training and experimental data, and the laboratory that verified the result. Existing law assigns ownership for inventions made by people. It is unprepared for discoveries assembled from delegated machine actions across organizations. Technical completion of a computation does not by itself create copyright, patent rights or property in its output.

This is not a remote legal puzzle. It determines whether the gains of a quantum-enabled science are broadly shared or narrowly captured. In a mature hybrid transformation, value is likely to concentrate at one or more of four control points:

- **The device layer.** Whoever builds and operates scarce quantum processors.
- **The cloud layer.** Whoever controls access, scheduling and pricing.
- **The data layer.** Whoever holds the experimental and proprietary data that make computation useful.
- **The governance layer.** Whoever defines the mandates, standards and verification that make results trustworthy and tradeable.

The balance among these control points is not determined by physics. It is a matter of institutional design, and it is still open. Quantum Civilization's research on ownership aims to make the alternatives explicit before the defaults harden.

CHAPTER 11

QCTM: The Quantum Civilization Transition Model

A causal architecture, not a score

The Quantum Civilization Transition Model is the institution's instrument for connecting capability to consequence without skipping the layers in between. It consists of four components:

1. **A versioned register of claims.** Every assertion about quantum progress and its effects, classified by evidence type.
2. **A causal map.** The hypothesized links between the layers of the Civilizational Compute Stack.
3. **An indicator dictionary.** Precisely defined measurements for each layer, with units, conditions, dates and sources.
4. **Conditional scenarios.** Alternative futures, each with its dependencies, turning points and refutation signals.

QCTM deliberately publishes **no composite civilization index**. Units, levels of measurement and causal pathways differ too much across layers, and there is no independently validated basis for weighting them. A single score would convey precision the evidence cannot support. Instead, every claim must travel the chain explicitly:

physical capability → logical computation → workflow time, cost and quality → adoption → scientific or economic outcome → institutional and social consequence

Each arrow is a hypothesis, not a law. The same outcomes are shaped by progress in classical computing, access to data, energy and labour costs, capital, policy, competition, skills, contracts and

the quality of measurement. Feedback is real. AI improves quantum calibration, and quantum results may in turn inform AI-driven science. Such cycles must be modelled as time-ordered stages, not as simultaneous correlations.




FIGURE 4 The QCTM causal chain. Every arrow is a hypothesis with its own test; feedback is modelled as time-ordered stages.

Indicators by layer

TABLE 4 The QCTM indicator dictionary (summary).

Layer	Core indicators	Rule
Physical device	Manufacturer, technology, device and date; qubit count; gate fidelity distributions; connectivity; readout time; calibration age; availability	Never compare qubit counts across technologies without workload and error profile [2]
Logical computation	Code type and distance; physical qubits per logical qubit; logical error per cycle; logical gate error and duration; decoding latency	Memory error is not gate error; preserve measurement conditions [1]
Hybrid workflow	Problem class and size; preprocessing; compilation; queue time; shots; retries; post-processing; error against a strong classical baseline	Report wall-clock and active time; include failed and cancelled jobs [2][13][14]
Access and cost	Price per task, shot, processor-second or reservation; currency, region, date; storage and transfer costs	A public price is not a contract price [26][28]
Agency and institution	Authorized vs. executed actions; human approval points; mandate duration; rejected jobs; reservation-invoice gap; evidence completeness	Do not measure “general intelligence” by task success rates
Science and economy	Pre-registered accuracy; independent replication; total human and machine cost per result; actual users and transactions	Counts of patents, papers or capital alone do not prove value
Security	Algorithm inventory; key sizes; data confidentiality horizon; migration status	Jurisdiction-specific; never inferred from device location alone [29][30]

Every QCTM record carries the indicator, its value and unit, the device and task, the measurement and publication dates, the source, the experimental conditions, the evidence type, the uncertainty, the reason for any missing data, the data publisher and its correction history.

The first entries

The founding edition of QCTM already contains documented observations. They are not rows in a common unit and imply no causal relation among themselves. Their purpose is to establish a verifiable starting line:


- **Quantum memory, 2024.** A distance-7 surface-code memory with $0.143\% \pm 0.003\%$ logical error per cycle on 101 physical qubits [1].
- **Logical memory fidelity, 2026.** 99.9995% reported on Helios using a colour code [51].
- **Fault-tolerance ingredients, 2026.** About forty-fold error reduction in cultivated magic states on Willow [52].
- **Shared benchmarks, 2026.** Ten optimization problem classes with open data and code in QOBLIB [14][15].
- **Access, 2026.** Up to ten minutes of quantum processor time per rolling 28 days on a public free plan [28]. Task-, shot- and reservation-based pricing on a multi-provider cloud [26].
- **Authorization, 2026.** A reported lifecycle-bound authorization protocol for quantum jobs, evaluated with simulators and synthetic inputs; publication record pending verification [22].
- **Autonomous science, 2023/2026.** Automated materials synthesis, with a subsequent published correction recorded as an integrity indicator [18][19].
- **Security, 2024/2026.** Three post-quantum standards approved, and the transition guidance finalized [29][30].

Causal tests and refutation

QCTM's strong claims are framed so that they can fail:

- **T1 · Workflow benefit.** At equal problem, solution quality and budget, a workflow with a quantum component reduces total time or cost relative to a strong classical baseline. If the benefit disappears against a better classical method, no quantum advantage is claimed for that task.
- **T2 · Mandate integrity.** A shared budget with atomic reservations reduces predefined cases of concurrent overspending relative to otherwise identical lifecycle-based control. If a strong general-purpose ledger performs as well, the new mechanism adds no value. If the difference appears only when a workload is *labelled* quantum, it is not quantum-specific.
- **T3 · Discovery effect.** With identical data and a pre-registered experimental budget, a research agent with a quantum component yields more independently verified findings than a classical agent. Claims about materials or synthesis require physical experiments. Computational candidacy alone proves nothing.
- **T4 · Institutional change.** An observable change in contractual or liability structures follows a specific use of quantum computation and is not better explained by AI, security requirements or cloud costs.

At model level, if successive pre-registered comparisons show no improvement on any predefined outcome, or if institutional effects are equally well explained by classical changes alone, QCTM's strong transition claims remain unconfirmed. Where data do not yet exist, a hypothesis is recorded as *untested*, never as refuted.



PART IV

The Research Programme

CHAPTER 12

Eight Research Frontiers

A research institution is defined less by its ambitions than by its choice of first questions. To choose them, Quantum Civilization assessed eight candidate research frontiers against three criteria. **Distinctiveness (D)** measures how clearly the contribution would differ from existing work. **Testability (T)** asks whether a small team can produce decisive evidence. **Magnitude (M)** is the scientific and practical value if the work succeeds. Scores from 1 to 5 are the institution’s initial expert judgements, to be refined by deeper literature review. They are not bibliometric measurements.

TABLE 5 Assessment of the eight research frontiers.

Frontier	D	T	M	Core question
A · Quantum Civilization Studies	2	2	4	Through which mechanisms does specialized computation reshape social power?
B · Quantum-Enabled Agent Economy	3	4	4	How do agents choose among CPU, GPU and QPU under budget and quality constraints?
C · Quantum Mandate Engineering	3	5	5	How can delegated, budgeted computation stay governable across agents and providers?
D · Hybrid Intelligence Infrastructure	2	3	4	Where can a governance-layer experiment improve AI-HPC-QPU integration?
E · Quantum Economic Transition	3	3	5	When does a workflow become cheaper or better, and to whom does the value flow?
F · Quantum Civilization Transition Model	3	4	4	Can capability-to-consequence claims be made explicit, measurable and refutable?
G · Verifiable Compute Delegation Across Providers	4	5	5	Does mandate-bound, multi-provider budget control prevent overspending better than strong existing solutions?
H · Quantum-Autonomous Science Integrity	3	3	4	Can an agent’s chain of evidence be reconstructed by an independent reviewer?

The assessment is candid about where the field is crowded. Hybrid AI-HPC-QPU orchestration is already a substantial engineering domain, with work on agent tool protocols, cloud orchestrators and fidelity-aware schedulers [21][23][24]. The general combination of “agents” and “quantum” is not new: Oak Ridge has studied multi-agent control planes for quantum networks [42], and theoretical work on quantum-enhanced agents exists [45]. Policy literature on quantum governance is growing [31][32][46]. None of this diminishes the programme. It defines precisely where a new institution can add what does not yet exist.

The decision. The first line of research combines G and C: *verifiable delegation of scarce computation across agents, organizations and providers*, tested with a fully classical workflow as a control. The second line combines E and F: the economic transition and the QCTM framework that measures it. Frontier A, together with the institution’s long-form book project, forms the public intellectual horizon. D and H are pursued in collaboration as partners emerge. If the first experiment shows that a well-implemented conventional ledger solves the problem equally well, that will be a genuine discovery. The quantum-specific claim will be narrowed, and the next testable question will be pursued.

CHAPTER 13

The First Experimental Programme

The five experiments below form the institution's first programme. Each is a proposal at protocol stage, not a completed study. Synthetic data, open datasets and any paid services will be labelled explicitly. Every experiment is designed so that a negative result is publishable and informative.

Experiment 1 · Shared budgets across agents and providers

Question. Does a shared, atomic reservation mechanism reduce unauthorized aggregate spending by concurrent agents, compared with local or per-job budget control? Does a new mechanism improve on Q-LEASE combined with a strong general-purpose ledger?

Design. Four conditions under identical identities, job phases, pricing and queue incentives: (1) provider API permissions with local per-job limits; (2) a faithful reconstruction of the reported Q-LEASE protocol [22], once its publication record is verified; (3) that protocol combined with a strong central atomic ledger; (4) the proposed mechanism. Agent counts of 1, 2, 10 and 50, with retries, duplicate requests, cancellations, delayed provider responses, reservation expiry, fixed exchange rates and price changes. Pricing structures are anchored in public documentation [26][28]. They are not claims about real customer logs.

Measurement. At every event, verify the invariant *confirmed cost + valid reservations* $\leq$ *budget*. Report overspend frequency and amount, false refusals, p50/p95 processing latency, throughput and audit completeness. Use at least 30 pre-registered seeds per scenario and method, Wilson intervals for binary rates and bootstrap intervals for latencies.

Quantum-specificity test. Run the same workflows with a QPU profile and with an ordinary GPU/CPU profile. A mechanism that helps only when a workload is labelled "quantum" has not demonstrated anything quantum-specific.

Outputs. Open simulator, event-trace schema, result tables including null results, and a manuscript. If condition 3 matches condition 4, no superiority is claimed.

Experiment 2 · Mandate-constrained hybrid scheduling

Question. What trade-offs does a mandate impose when one set of jobs chooses among CPU, GPU and QPU with different prices, fidelities and queues? The hypothesis is that mandate control prevents impermissible decisions at a measurable cost in waiting time and throughput. No claim is made in advance that every objective improves.

Design. Three synthetic provider profiles whose units and pricing structures follow public documentation. Compare FIFO, a fair-share approximation derived from published principles [25], cheapest-provider selection and, where available, the public QFOR implementation [23]. Then add the same mandate control to every method. Any approximation of a proprietary scheduler is labelled as such.

Measurement. Accepted-job share, p50/p95 completion time, actual cost, result quality, waiting-time distribution for the smallest users and access by organization, across at least 30 independent seeds.

Experiment 3 · Whole-workflow cost comparison

Question. For small, pre-fixed problems, does a hybrid quantum workflow differ from a classical one in solution quality, time and cost? The hypothesis is two-sided. On small hardware the classical solution may well prevail, and the experiment is designed to show this cleanly if so.

Design. Suitable small instances from QOBLIB [14][15] and public Max-Cut or QUBO sets with documented parameters. Exact classical solutions where feasible and strong public heuristics, compared with simulated QAOA workflows in standard toolkits. Benchmark circuits from MQT Bench [50]. An optional micro-check on real hardware runs only within free quotas or an approved budget [28].

Measurement. Full wall-clock time, queue, compute time, all repetitions, classical optimizer cost, distance from known optima, monetary cost and software versions. Energy is reported only if actually measured. A simulator cannot demonstrate quantum speed-up, and this experiment makes no such claim.

Experiment 4 · The public QCTM register

Question. Does a layered register of claims and indicators produce less contested and more verifiable expert judgements than a single success score?

Design. Extend the founding entries (Chapter 11) to 10–20 sources. At least two independent assessors blindly classify each entry's status and missing assumptions, separating source fact, interpretation and testable hypothesis. The agreement metric is registered before assessment. Consent and, where required, ethical review are secured for human participants.

Outputs. JSON/CSV register, methodology note and expert commentary published anonymously or with consent.

Experiment 5 · The evidence chain of a research agent

Question. Can an independent reviewer reconstruct which hypothesis an agent tested, with which data, model, tools and budget, and whether its output was actually verified? The hypothesis is that a structured evidence chain improves reproducibility over free-text logs.

Design. For ten small synthetic or public tasks, produce a free-text log and a machine-readable event record. Two experts attempt blind reconstruction of source data, exact versions, authorization and result. Compare success rate, time and falsely confirmed results. The first phase may use simulators. No new substance is claimed as discovered [17][18][21].

Outputs. A reconstruction protocol for autonomous science. If reviewers cannot independently re-run a trace, the manuscript is marked incomplete.

CHAPTER 14


The Standard We Hold Ourselves To

An institution that studies how autonomous systems should be governed must govern its own research with equal rigour. Quantum Civilization adopts the following standards from its founding:

1. **Pre-registration.** Research questions, protocols, data sources, strong baselines and primary outcome measures are published, where possible, before results are seen.
2. **Versioning.** Every result carries a version and a date. Code, pinned dependencies, run instructions, data dictionaries and limitations are published to the extent licences and data protection allow.
3. **Complete reporting.** Null results, unfavourable results and results that undermine any company's claims are preserved. Simulations, hardware experiments, company announcements and peer-reviewed work are labelled as distinct evidence classes.

4. **Correction.** A published procedure governs corrections and retractions. Scientific editors and external reviewers may challenge conclusions, including those the founder prefers.
5. **Independence and disclosure.** Funding sources, ownership relations and overlapping roles are disclosed. No funder or founder may suppress an unwelcome result or change a method after the fact without a public correction notice.
6. **Ethics and data protection.** Assessments precede any work involving people, financial data, employees or confidential logs. Demonstrations never use real payments or personal data.
7. **Safety boundaries.** Success in a research simulator is never presented as sufficient for controlling safety-critical production systems.

These standards are not a disclaimer appended to an ambitious programme. They are what makes the programme credible. An institution earns the right to ask large questions by showing in public how it answers small ones.



PART V

Horizons

CHAPTER 15

Beyond the Classical World: Three Futures, 2027–2042

Scenarios are not forecasts. They are disciplined descriptions of internally consistent futures, each defined by the mechanisms that would produce it, the signposts that would announce it and the evidence that would refute it. Their purpose is not to predict which future will arrive. It is to recognize early which one is arriving. Quantum Civilization tracks three.




FIGURE 5 Three futures, 2027–2042. Paths diverge as signposts accumulate; each carries its own refutation signal.

Scenario I • The classical future still happens

Mechanism. Artificial intelligence and automation transform science, industry and government on classical hardware. Fault-tolerant quantum computation progresses, but its cost, overheads and access barriers keep it confined to specialized scientific and security applications. Classical algorithms continue to absorb claimed quantum advantages.

What the world looks like. The quantum age arrives first, and for a long time mainly, as a cryptographic migration. Power concentrates where classical AI is built and run: in data centres, energy supply, chip fabrication and model ownership. The institutional questions of Mandate Engineering become urgent anyway, because autonomous agents spend classical compute on a massive scale.

Signposts. Repeated dequantization of proposed advantages; logical-gate costs that fall more slowly than roadmaps; quantum cloud usage that stays concentrated in research.

Refutation signal. A verified, economically relevant workflow in which a quantum component beats the strongest classical baseline on total cost.

Scenario II • Quantum advantage arrives in narrow channels

Mechanism. Selected chemistry, materials and optimization workflows achieve verified advantage on fault-tolerant or early fault-tolerant machines. Autonomous agents begin to purchase quantum computation as a premium service within larger classical pipelines.

What the world looks like. Quantum computation becomes a strategic input to a limited number of industries and national programmes. The decisive questions become contractual: who may commis-

sion the job, at what price, under which mandate, and who owns the result. Workflow-level contracts, verification services and specialized brokers emerge.

Signposts. Published end-to-end advantage on industrially relevant chemistry or materials problems; pricing models built around logical rather than physical operations; agent frameworks that natively route work to quantum providers.

Refutation signal. Advantage claims that fail independent replication or collapse when total workflow cost is counted.

Scenario III · The hybrid transformation

Mechanism. Error correction, algorithms and infrastructure mature together. For whole classes of problems the cost of simulation-driven discovery collapses, and autonomous research systems integrate quantum computation as routinely as they now use GPUs.

What the world looks like. Discovery in chemistry, materials and parts of biology accelerates measurably. Economic power gathers at whichever control points dominate: devices, clouds, data or governance. Computational sovereignty becomes a first-order concern of states. Whether the gains are broadly shared or narrowly captured becomes the defining political question of the era. Technical success alone does not settle it.

Signposts. Logical-qubit counts and logical-gate fidelities on credible exponential trajectories; autonomous laboratories publishing quantum-assisted, independently replicated discoveries; the first legal disputes over ownership of machine-commissioned results.

Refutation signal. Persistent failure of logical computation to scale economically despite continued investment.

A speculative horizon beyond 2042

Beyond the scenario period, the institution maintains one explicitly speculative horizon, kept separate from its evidence-based work. It is a civilization in which the most consequential decisions about matter, energy and biology are informed by simulations no human could perform and no classical machine could complete, commissioned by autonomous systems acting under mandates that encode human purposes. In such a world the central question of political philosophy, *who decides?*, acquires a computational form: *who authorizes the machines that know?* Quantum Civilization does not predict this future. It studies the decisions being made now that would make it more just, more accountable and more widely shared if it arrives.

CHAPTER 16

The Questions That Define the Territory

The founding thesis generates five questions. For each, the institution states what is known, what remains unresolved and why the answer matters.

1. What happens when autonomous intelligence, quantum computation, robotics, scientific discovery and synthetic economic systems become interconnected?

Known. Each component exists and is advancing. Agents operate laboratory instruments and submit quantum jobs through standard protocols [17][18][21]. Error-corrected quantum memory scales below threshold [1][51].

Unresolved. Whether the interconnected system exhibits behaviour — speed, failure modes, concentrations of control — that its components do not.

Why it matters. Interconnected infrastructures fail and concentrate power in ways their parts do not. The time to study the system is before it is built.

2. How might the architecture of scientific discovery change?

Known. Autonomous laboratories can plan and execute experiments, and their errors can propagate until verification catches them [18][19][20].

Unresolved. Whether quantum simulation, used by autonomous agents, measurably increases the rate of independently verified discovery (test T3).

Why it matters. If discovery becomes a metered, delegated process, the rules of authorization and verification become the rules of science itself.

3. Who controls computational resources?

Known. Quantum access is priced, queued and permissioned by a small number of providers and national programmes [25][26][28][31].

Unresolved. Whether control will concentrate at the device, cloud, data or governance layer, and whether public infrastructure can provide equitable access.

Why it matters. In a computational civilization, control of computation is a form of sovereignty.

4. Who owns the value generated by autonomous systems?

Known. Technical completion of a computation does not itself create ownership. Existing law assigns inventions to people, not to delegated machine actions.

Unresolved. How ownership will be allocated among principals, model providers, compute operators, data holders and verifying laboratories.

Why it matters. The answer determines whether the gains of a quantum-enabled science are shared or captured.

5. How will institutional power, economic sovereignty and human agency evolve?

Known. Institutions are already reorganizing around anticipated quantum capability [29][30][46].

Unresolved. Whether humans retain meaningful authority over autonomous systems that commission computation beyond human comprehension, and through which mechanisms.

Why it matters. Human agency in a computational civilization will be exercised largely through mandates, and mandates can be well or badly engineered.

EPILOGUE

Earning the Ambition

Quantum Civilization begins with a large hypothesis and small, exact experiments. That combination is deliberate. Ambition sets the direction of an institution. Evidence alone determines how far it may claim to have travelled.

The path is defined in stages, each closed by a decision gate:

- **The first thirty days.** Fix the core problem and the concepts of Mandate Engineering. Complete an in-depth review of Q-LEASE, QFOR, Qonductor and multi-party agent accounting. Publish a data dictionary, a threat model and a draft pre-registered protocol. *Gate: are the strong baselines and the originality question stated precisely?*
- **The first ninety days.** Build the minimal simulator for Experiment 1, collect pre-registered results and publish favourable and unfavourable findings with code. Extend the QCTM register and obtain two independent written reviews. *Gate: does any result survive comparison with a strong general-purpose baseline?*
- **The first year.** At least one open working paper and one manuscript submitted for external peer review; a reproducible benchmark release; a sample chapter of the book.
- **The first three years.** A network of independent partner institutions, several reproducible public studies, and a benchmark or method recognized by others. If the quantum-specific distinction does not hold, the programme continues as a science of delegated hybrid computation.

The institution will be judged by others' independent use, replication and criticism of its work, not by the boldness of its statements. That is the standard it sets for itself. The questions it asks are among the most consequential of the century, and they deserve answers that can withstand scrutiny.

APPENDIX

Method and Evidence Standard

Scope. This paper rests on a targeted mapping of literature and institutions as of 28 October 2026. It is not a PRISMA-compliant systematic review and does not claim to list every laboratory in the world. The search covered peer-reviewed articles, preprints, governmental and institutional sources, open-source software and provider documentation. It included primary sources in Chinese, Japanese and Korean. Absence of a term from search results does not prove the absence of a research niche.

Evidence classes. The paper distinguishes: *peer-reviewed results* (published in journals or conference proceedings); *preprints* (not equivalent to peer review); *company announcements and technical documentation* (evidence of a stated plan, price or service property, not of an independent breakthrough); *roadmaps* (engineering intentions); *projections* (claims about the future); *scenarios* (conditional constructions, not probability statements); and *hypotheses* (claims designed to be tested).

Limits. Microdata on quantum cloud queues and confidential contract prices are not publicly available. This edition therefore makes no market-size estimate, bibliometric census or independent device benchmark. Research frontier scores are initial expert judgements. One source, Q-LEASE [22],

was identified in the source mapping, but its publication record could not be independently confirmed at the time of writing. It is treated as a candidate baseline that must be verified before Experiment 1 begins; if it cannot be, the experiment proceeds with the remaining baselines.

Authorship and preparation. Author: Steven Alber. Publisher: Quantum Civilization. Literature mapping and drafting were supported by AI research tools working under the author's direction. Every factual claim is tied to the cited sources, and responsibility for the text rests with the author.

REFERENCES


Sources

- 1 Google Quantum AI and collaborators. Quantum error correction below the surface code threshold. Nature (2024).
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-024-08449-y>
- 2 Proctor, T., Young, K., Baczewski, A. D. & Blume-Kohout, R. Benchmarking quantum computers. Nature Reviews Physics (2025).
PEER-REVIEWED REVIEW
<https://www.nature.com/articles/s42254-024-00796-z>
- 3 Google Quantum AI and collaborators. Observation of constructive interference at the edge of quantum ergodicity. Nature, 22 October 2025.
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-025-09526-6>
- 4 Kim, Y. et al. Evidence for the utility of quantum computing before fault tolerance. Nature, 14 June 2023.
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-023-06096-3>
- 5 Bluvstein, D. et al. Logical quantum processor based on reconfigurable atom arrays. Nature (2023).
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-023-06927-3>
- 6 PsiQuantum team. A manufacturable platform for photonic quantum computing. Nature, 26 February 2025.
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-025-08820-7>
- 7 Xanadu and collaborators. Scaling and networking a modular photonic quantum computer. Nature, 22 January 2025.
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-024-08406-9>
- 8 Legg, H. F. On the robustness of topological gap detection via transport. Nature (2026).
PEER-REVIEWED ANALYSIS
<https://www.nature.com/articles/s41586-026-10567-8>
- 9 Artificial intelligence for quantum computing. Nature Communications (2025).
PEER-REVIEWED REVIEW
<https://www.nature.com/articles/s41467-025-65836-3>
- 10 Reinforcement learning control of quantum error correction. Nature, 8 July 2026.
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-026-10759-2>
- 11 Huang, H.-Y. et al. Power of data in quantum machine learning. Nature Communications (2021).
PEER-REVIEWED STUDY
<https://www.nature.com/articles/s41467-021-22539-9>
- 12 Tang, E. A quantum-inspired classical algorithm for recommendation systems. arXiv:1807.04271 (2018).
PREPRINT (CITED VERSION)
<https://arxiv.org/abs/1807.04271>
- 13 Quantum annealing applications, challenges and limitations for optimisation problems compared to classical solvers. Scientific Reports, 13 April 2025.
PEER-REVIEWED COMPARISON
<https://www.nature.com/articles/s41598-025-96220-2>

- 14 Koch, T. et al. The Quantum Optimization Benchmarking Library. Nature Computational Science, 23 June 2026.
PEER-REVIEWED ARTICLE
<https://pmc.ncbi.nlm.nih.gov/articles/PMC13293884/>
- 15 QOBLIB — Quantum Optimization Benchmarking Library. Open benchmark code and data. Accessed 28 October 2026.
OPEN DATASET AND SOFTWARE
<https://github.com/ZIB-AOPT/QOBLIB>
- 16 Option pricing under stochastic volatility on a quantum computer. Quantum 8, 1504, 23 October 2024.
PEER-REVIEWED RESOURCE ESTIMATE
<https://doi.org/10.22331/q-2024-10-23-1504>
- 17 Boiko, D. A. et al. Autonomous chemical research with large language models. Nature, 20 December 2023.
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-023-06792-0>
- 18 Szymanski, N. J. et al. An autonomous laboratory for the accelerated synthesis of inorganic materials. Nature, 29 November 2023 (corrected version).
PEER-REVIEWED EXPERIMENT
<https://www.nature.com/articles/s41586-023-06734-w>
- 19 Szymanski, N. J. et al. Author Correction: An autonomous laboratory for the accelerated synthesis of inorganic materials. Nature, 19 January 2026.
AUTHOR CORRECTION
<https://www.nature.com/articles/s41586-025-09992-y>
- 20 Panapitiya, G. et al. AutoLabs: cognitive multi-agent systems with self-correction for autonomous chemical experimentation. Scientific Reports, 25 June 2026.
PEER-REVIEWED STUDY
<https://www.nature.com/articles/s41598-026-45593-z>
- 21 Shiraishi, Hamamura, Ishigaki & Kadowaki. A Model Context Protocol Server for Quantum Execution in Hybrid Quantum-HPC Environments. IEEE QCNC 2026; arXiv preprint.
CONFERENCE PAPER AND PREPRINT
<https://ieeexplore.ieee.org/document/11501229/>
<https://arxiv.org/abs/2604.08318>
- 22 Ailabouni & Román-Gallego. Q-LEASE: Lifecycle-bound zero-trust authorization for HPC-to-QPU job handoffs. Open Research Europe, 24 August 2026.
REPORTED PUBLICATION; RECORD NOT INDEPENDENTLY CONFIRMED AT TIME OF WRITING
<https://ore.eu/ore/article/view/1497>
- 23 Nguyen, Usman & Buyya. QFOR: A Fidelity-aware Orchestrator for Quantum Computing Environments using Deep Reinforcement Learning. arXiv, 7 August 2025.
PREPRINT (CITED VERSION)
<https://arxiv.org/abs/2508.04974>
- 24 Qonductor: A Cloud Orchestrator for Quantum Computing. SC 2025.
CONFERENCE PAPER (PUBLIC MANUSCRIPT)
<https://dse.in.tum.de/wp-content/uploads/2025/09/Qonductor-SC-2025.pdf>
- 25 IBM Quantum Documentation. Fair-share scheduler. Accessed 28 October 2026.
TECHNICAL DOCUMENTATION
<https://quantum.cloud.ibm.com/docs/en/guides/fair-share-scheduler>
- 26 Amazon Web Services. Amazon Braket pricing. Accessed 28 October 2026; prices and terms change.
PUBLIC PRICING DOCUMENTATION
<https://aws.amazon.com/braket/pricing/>
- 27 Amazon Web Services. Actions, resources, and condition keys for Amazon Braket. Accessed 28 October 2026.
OFFICIAL IAM DOCUMENTATION
https://docs.aws.amazon.com/service-authorization/latest/reference/list_braket.html
- 28 IBM Quantum Documentation. Plans overview. Accessed 28 October 2026; terms may change.
TECHNICAL DOCUMENTATION
<https://quantum.cloud.ibm.com/docs/en/guides/plans-overview>
- 29 NIST. Announcing approval of three Federal Information Processing Standards for post-quantum cryptography. 13 August 2024.
OFFICIAL STANDARDS ANNOUNCEMENT
<https://www.nist.gov/news-events/news/2024/08/announcing-approval-three-federal-information-processing-standards-fips>
- 30 NIST. IR 8547: Transition to Post-Quantum Cryptography Standards (final). 27 October 2026.
OFFICIAL TECHNICAL REPORT
<https://csrc.nist.gov/pubs/ir/8547/final>
- 31 OECD. An overview of national strategies and policies for quantum technologies. 2025.
OFFICIAL POLICY STUDY
https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/12/an-overview-of-national-strategies-and-policies-for-quantum-technologies_33a0b249/5e55e7ab-en.pdf

- 32 UNESCO. The Quantum Moment: A Global Report. 2026.
INSTITUTIONAL REPORT
<https://www.unesco.org/en/articles/quantum-moment-global-report>
- 33 Gao, D. et al. Establishing a new benchmark in quantum computational advantage with 105-qubit Zuchongzhi 3.0 processor. Physical Review Letters 134, 090601, 7 March 2025.
PEER-REVIEWED EXPERIMENT
<https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.134.090601>
- 34 Chinese Academy of Sciences. Zuchongzhi 3.0 announcement (in Chinese). 5 March 2025.
INSTITUTIONAL ANNOUNCEMENT
https://bfse.cas.cn/yjz/202503/t20250305_5048759.html
- 35 RIKEN. Integration of a quantum computer with the Fugaku supercomputer in quantum chemistry, 19 June 2025; RIKEN and Fujitsu 256-qubit superconducting system, 22 April 2025 (in Japanese).
OFFICIAL INSTITUTIONAL SOURCES
https://www.riken.jp/press/2025/20250619_1/
https://www.riken.jp/pr/news/2025/20250422_1/index.html
- 36 KAIST Graduate School of Quantum Science and Technology; KAIST-KRISS superconducting systems programme. Accessed 28 October 2026.
OFFICIAL INSTITUTIONAL PAGES
<https://quantum.kaist.ac.kr/main>
<https://quantum.kaist.ac.kr/english/research-kriss>
- 37 Centre for Quantum Technologies, Singapore. National role; National Quantum Office and Quantinuum strategic partnership, 6 November 2025.
OFFICIAL INSTITUTIONAL SOURCES
<https://www.cqt.sg/national-role/>
<https://www.cqt.sg/highlight/2025-11-nqo-quantinuum-strategic-partnership/>
- 38 Government of India, Department of Science and Technology. National Quantum Mission, 2023–2031.
OFFICIAL PROGRAMME DESCRIPTION
<https://dst.gov.in/national-quantum-mission-nqm>
- 39 UNSW Sydney. Quantum computer chips clear major manufacturing hurdle. 25 September 2025.
UNIVERSITY ANNOUNCEMENT
<https://www.unsw.edu.au/newsroom/news/2025/09/quantum-computer-chips-clear-major-manufacturing-hurdle>
- 40 UK National Quantum Computing Centre. Quantum computing testbeds. Accessed 28 October 2026.
OFFICIAL PROGRAMME DESCRIPTION
<https://www.nqcc.ac.uk/quantum-computing-testbeds-in-the-uk/>
- 41 European Commission. Quantum Europe strategy and quantum infrastructure, 2025.
OFFICIAL POLICY DOCUMENTS
<https://digital-strategy.ec.europa.eu/en/policies/quantum>
<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52025DC0363>
- 42 Alnajjar, Zhang & Kiran. Multi-Agent Control Planes for Quantum Networks. Oak Ridge National Laboratory; GLSVLSI, June 2026.
CONFERENCE PAPER
<https://www.ornl.gov/publications/multi-agent-control-planes-for-quantum-networks-a-scalable-architecture-for-autonomous-quantum-internet-management/>
- 43 Quantinuum and Microsoft. Demonstration of reliable logical qubits. 3 April 2024.
COMPANY ANNOUNCEMENT
<https://www.quantinuum.com/press-releases/quantinuum-and-microsoft-announce-new-era-in-quantum-computing-with-breakthrough-demonstration-of-reliable-qubits>
- 44 IonQ. Publications. Accessed 28 October 2026.
COMPANY RESEARCH CATALOGUE
<https://www.ionq.com/publications>
- 45 Quantum and Complexity Science Initiative. Research fellowships in quantum-enhanced agents, learning and algorithms. Accessed 28 October 2026.
RESEARCH PROGRAMME DESCRIPTION
<https://www.quantumcomplexity.org/join-us/research-fellowships/>
- 46 OECD. Recommendation of the Council on Quantum Technologies, OECD/LEGAL/0508. Adopted 28 May 2026.
INTERNATIONAL POLICY INSTRUMENT
<https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0508>
- 47 Amazon Web Services. Amazon Braket quotas. Accessed 28 October 2026.
TECHNICAL DOCUMENTATION
<https://docs.aws.amazon.com/braket/latest/developerguide/braket-quotas.html>
- 48 University of Waterloo, Institute for Quantum Computing. Quantum information with trapped ions: processors. Accessed 28 October 2026.
RESEARCH GROUP DESCRIPTION
<https://qiti.iqc.uwaterloo.ca/processors/>

- 49 IBM. Quantum development and innovation roadmap. Accessed 28 October 2026.
COMPANY ROADMAP (PLANS, NOT RESULTS)
<https://www.ibm.com/roadmaps/quantum/>
- 50 Munich Quantum Toolkit. MQT Bench. Open benchmark software. Accessed 28 October 2026.
OPEN SOFTWARE
<https://github.com/munich-quantum-toolkit/bench>
- 51 Quantinuum. Quantinuum breaks error correction world record by 10x, validating its leading path to fault-tolerant quantum computing by 2029. 19 October 2026.
COMPANY-REPORTED EXPERIMENT
<https://www.quantinuum.com/press-releases/quantinuum-breaks-error-correction-world-record-by-10x-validating-its-leading-path-to-fault-tolerant-quantum-computing-by-2029>
<https://thequantuminsider.com/2026/10/19/quantinuum-reports-a-10x-reduction-in-logical-error-rate-on-helios/>
- 52 Google Quantum AI. Magic state cultivation on the Willow processor. Science (2026); Google announcement, 15 October 2026.
PEER-REVIEWED EXPERIMENT
<https://doi.org/10.1126/science.aee7961>
<https://blog.google/innovation-and-ai/models-and-research/quantum-computing/magic-state-cultivation-quantum-computing/>



QUANTUM
CIVILIZATION
RESEARCH INSTITUTE

Physics will decide what becomes computable.
Institutions will decide what becomes of it.