



QubitCore

Space Qualification of a Complete Entanglement Distribution Payload

Michal Krelina, Alper Ozulker, Zeki Seskir, Utku Tefek, Kadir Durak
Qubitrium




Table of Contents

Summary for Non-technical Readers.....	2
1. Summary.....	2
2. Why Space Entanglement Distribution Is Needed.....	5
2.1 What entanglement distribution is for.....	5
2.1.1 Keys.....	6
2.1.2 Time and frequency transfer.....	6
2.1.3 Telescopes and sensors working as one instrument.....	6
2.1.4 Distributed quantum computation.....	6
2.2 Why QKD comes first.....	6
2.3 Distance, and why relays are the real problem.....	7
2.4 What has been demonstrated.....	8
2.5 The bottleneck and the market behind it.....	9
3. QubitCore.....	10
3.1 Why the payload runs the full protocol on board.....	10
3.2 Architecture.....	11
3.3 What flew, and what it qualifies.....	12
4. Differentiation.....	13
4.1 A complete entanglement link in 1U.....	13
4.2 A product, not a mission.....	14
4.3 Heritage.....	14
5. Roadmap: From Qualified Hardware to the Quantum Internet.....	15
5.1 The network Qubitrium is building toward.....	15
5.2 The steps from here.....	16
6. Conclusion.....	18
7. References.....	19

Summary for Non-technical Readers

Governments, banks and other organisations send messages that must stay secret for many years. Today, these messages are protected by encryption based on mathematics. Future quantum computers may break some of this encryption, and new mathematical methods are the main answer.

For a few very sensitive links, a second layer of protection is possible. It uses single particles of light to create secret keys, the numbers that lock and unlock encrypted messages. Anyone who tries to copy this light leaves a trace that both sides can detect. The same light can also share very precise time between distant places.

This light fades in glass fibre, and it cannot be boosted along the way. Over more than a few hundred kilometres, it has to come from a satellite. A Chinese satellite showed this in 2017. Since then, the problem has been the equipment. Nobody sells it ready to fly, so every mission builds its own, and this makes missions slow and expensive.

Qubitrium has built this equipment as a product. It fits in a 10 cm cube and weighs less than one kilogram. It was launched into space in March 2026, and in orbit it has made secret keys on its own, with no help from the ground. On this first flight, the light stays inside the device, so nothing is sent to the ground yet.

Later versions add a telescope that sends the light to a ground station. After that, a second telescope is planned, so that one satellite serves two ground stations at once, up to about 2,000 km apart. These two stations can then share a secret key that even the satellite does not know.

Qubitrium reached space with a few per cent of the money behind the large government programmes. It sells these parts as products, so other missions can buy them instead of building their own.

1. Summary

A quantum internet links distant sites by sharing entangled photons. Entanglement is a correlation between two particles: measure one and you learn something about the other, and no third party can hold a copy. A quantum internet therefore does not carry data — it carries this correlation, and quantum applications are built on top of it. Today those applications are distributing secure encryption keys and precise time transfer. Soon, networked sensor arrays that act as one instrument, and functional quantum computers that share quantum states and compute will prevail. The list is non-exhaustive: the standard review of the field, written in 2018, expected more uses within the decade [16].

Over continental distances, entanglement has to travel through space. Optical fibre absorbs light, and nothing can amplify a quantum signal along the way, so past a few hundred kilometres only two options remain. One is a chain of relay stations, but every station holds the key in readable form, which means whoever runs a station can read the traffic passing through it. The other is a satellite. Quantum repeaters may become a third option later, but the technology is not ready, and this is why the major national programmes now include a space segment.

Physicists settled the science experimentally with the first satellite-to-ground results in 2017 [9, 10], and the obstacle since then has been the hardware. Nobody sells entangled photon sources, single-photon detectors and picosecond timing electronics as a qualified, ready-to-fly product, so every mission so far has built its own. That is a large part of why these missions are expensive. A small satellite demonstrator costs from a few million to a few tens of millions. A dedicated satellite costs around a hundred million and takes close to a decade.

We therefore qualified that hardware first, on its own, before joining it to an optical link.

QubitCore is the result: a 1U payload, a 10 cm cube weighing 950 g, holding a complete entanglement distribution system. That means the photon source, both measuring stations, the detectors, the timing electronics and the software that turns detections into keys. Both photons of each pair go to receivers inside the payload, so the payload runs the full protocol on board. In March 2026, Space-X Transporter-16 mission brought it into low Earth orbit, as a hosted payload sharing a ride on a D-Orbit carrier.

This first flight qualifies the hardware. Because the protocol runs on board from pair generation through to key output, the payload checks its own health in flight. We made it self-contained on purpose, because separating the quantum subsystems from the optical link lets us measure their performance in orbit on its own. From here the payload gains a telecom-wavelength source, then an optical terminal, then a second downlink, and finally quantum memory. The quantum payload itself changes little across those steps, which is why we qualified it first and separately.

QubitCore shows that the quantum core of an entanglement node can be small, cheap and space-worthy, and we reached orbit on a small fraction of the capital behind the institutional programmes we intend to supply.

Glossary

Entanglement. A correlation between two particles created together: measure one and you learn something about the other, and no third party can hold a copy.

Entangled photon pair. Two particles of light made at the same instant and entangled with each other. This is what a source makes and what a network distributes.

Quantum key distribution (QKD). Using single photons to give two distant sites the same secret number, which they then use as an encryption key. With entangled pairs, the key comes from the matching measurements at each end.

Trusted node. A relay station that holds the key in readable form, so whoever controls the station can read the traffic passing through it.

Single-photon detector. A device sensitive enough to register one particle of light and record when it arrived.

Quantum memory. A device that holds a photon's quantum state until it is needed, so links made at different times can be combined.

Quantum repeater. A station built from quantum memories that extends entanglement over long distances without ever holding the key.

Entanglement swapping. Interfere one photon from each of two separate pairs and the other two become entangled, although they never met. This is the basic step of a quantum repeater.




Figure 1: The QubitCore flight payload. A complete entanglement distribution system in 1U, 950 g. The boards on its faces are the single-photon detector channels; the coaxial cables carry their output pulses to the time tagger.

2. Why Space Entanglement Distribution Is Needed

2.1 What entanglement distribution is for

Two sites that share entanglement get matching measurement results that no ordinary channel can reproduce and no third party can copy. Everything a quantum network does rests on one or both of those facts. Coordination between distant sites, such as timing and sensing, comes from the first, and secrecy comes from the second, because a correlation nobody else can share is private by construction. The field's own model for planning the quantum internet rests on the same two facts [16].

Entanglement has four main uses. The first two work with the hardware described in this paper; the other two need quantum memory, which comes later.

2.1.1 Keys

Two ground stations each receive one photon of a pair, and from their matching measurements they build a shared key that nobody else holds, not even the operator of the satellite that made the pairs. That last part holds as long as both ground stations are themselves trustworthy and well built.

2.1.2 Time and frequency transfer

The two photons of a pair are born at the same instant, so the same detections that make a key can also transfer time. Our recent assessment of these protocols sets out what works today, and where published performance still falls short of theory [25].

2.1.3 Telescopes and sensors working as one instrument

The best developed proposal links distant telescopes into one instrument, with shared entanglement replacing the need to physically carry light from one telescope to the other [26]. Because it needs quantum memory at each telescope, it arrives after keys and timing.

2.1.4 Distributed quantum computation

A user with a simple quantum terminal can run a computation on a remote processor without the processor learning anything about it, and separate processors can link into one machine [16]. This too needs memory at the nodes, and it is the furthest away of the four.

Entanglement is a general-purpose resource, and we have not found all its uses yet [16]. A network that distributes entanglement serves the uses we know and the ones we do not; a network that distributes keys serves only keys.

2.2 Why QKD comes first

QKD reaches the market first, and it is much the easiest of these services to deliver. It needs no quantum memory and nothing beyond the pairs themselves and a detector at each end, and it is the one service with procurement programmes and buyers behind it today. Time transfer sits in the same category, since it uses the same pairs and the same detections and asks nothing extra of the payload. Together they carry the commercial case for the next payload.

Post-quantum cryptography is the main answer to the quantum threat against RSA and elliptic-curve cryptography, and it will carry the great majority of traffic [1, 2]. NIST

published the standards in August 2024 [4], and its draft schedule retires the old algorithms after 2030 and bans them after 2035 [3]. QKD is a partial and complementary measure, worth its cost on a small number of high-value links. It runs on top of conventional cryptography rather than replacing it, because the ordinary data link alongside the quantum one still has to be authenticated by conventional means. The reason to add it is defence in depth. Combine the two keys and an attacker has to break both layers, one of which rests on a hard mathematical problem and the other on the physical behaviour of light. What limits QKD is distance.

2.3 Distance, and why relays are the real problem

Loss in optical fibre grows exponentially with distance, and nothing can amplify a quantum signal along the way, because copying an unknown quantum state is impossible [5]. That puts a hard ceiling on how much key a fibre can carry [6]. The Micius team put the consequence plainly: over 1,200 km of fibre, even with a perfect transmitter running at ten billion pulses per second and flawless detectors with no noise, you would wait six million years for a single bit of key [9]. Twin-field QKD pushes past that ceiling and now spans 1,002 km in the laboratory [12], but it needs a measuring station in the middle and runs far too slowly for anyone to sell. That middle station does not have to be trusted, which is what separates it from the relay chains described below.

Terrestrial QKD therefore stays metropolitan, and anything longer is built from trusted relays. China's network chains more than 700 fibre links along a 2,000 km backbone, serves over 150 users, and reaches 4,600 km end to end once its satellite links are counted [14]. It works, but the limit is in its design rather than its build quality: every relay decrypts the key and encrypts it again, so the whole chain is only as secure as its weakest node. For a commercial operator that is a cost; for a government it is a sovereignty problem, because national keys then pass through equipment somebody else controls.

A satellite that makes the keys on board does not solve this problem, it moves it into orbit. The spacecraft generates the key, stores it, and hands it to each ground station in turn, so the operator knows every key. That is how the intercontinental results worked, including the recent China to South Africa exchange over more than 12,900 km [15].

Trusted-node satellites work, they deliver keys today, and they will stay the practical option until entangled links to two ground stations at once become routine. Our first downlink mission, the step after EPIQ (Section 5.2), sits in this category: an entangled source on the spacecraft, one photon of each pair measured on board and the other sent down, with the satellite still holding the key. Micius demonstrated exactly this in 2017 [10].

The important part is that the hardware does not change afterwards. Send the second photon to a second ground station instead of to an onboard detector, and the satellite

never sees the key: what remains is an entanglement link between two points on the ground, and keys are only one use for it. A satellite that generates keys on board cannot make that change, so you have to replace it.

The endpoint is a satellite that sends one photon of each pair to each of two ground stations and never holds the key at all, because no key exists until both parties measure.

This architecture has a geometric limit: both stations must be visible at once. From a 510 km orbit at 20° minimum elevation the footprint radius is about 1,060 km, so a single pass can serve stations up to roughly 2,100 km apart, or about 3,200 km at 10°. That covers continental pairs, and the 1,120 km demonstration [11] sits well inside it. Going further needs quantum memory on board, or a fall back to trusted-node operation.

Fibre also reaches only where somebody has laid it, which leaves out ships, aircraft, offshore platforms, deployed units and remote sites. Nearly all international capacity runs through submarine cables, with satellites carrying under one per cent [22], and those cables are easy to reach at their landing points and along their routes.

2.4 What has been demonstrated

Year	Result	Key figures	Ref.
2017	Satellite-to-ground QKD, keys made on board	1.1 kbit/s final secure key over a single 273 s pass, 1.1% error rate; 300 mm satellite and 1 m ground apertures; 22 dB diffraction loss at 1,200 km	[9]
2017	Entanglement-based QKD, one photon measured on board	3.5 bit/s average over 530 to 1,000 km; downlink loss 29 to 36 dB; satellite is a trusted node	[10]
2020	Entanglement-based QKD, both photons sent down, no trusted relay	0.12 bit/s between stations 1,120 km apart	[11]
2020	Entangled photon source on a nanosatellite	Bell test passed in orbit, 2.6 ± 0.06 , where any value above 2 confirms entanglement	[13]
2021	Integrated space-to-ground network	4,600 km, 47.8 kbit/s average per pass, over 150 users, trusted relays throughout	[14]
2025	Microsatellite QKD, real-time	22.7 kg payload, 1.07 Mbit of key per pass, China to South Africa over 12,900 km via trusted node	[15]

The 2025 result is the strongest outside support for our approach, and its authors state the problem directly: scaling to a constellation requires “small lightweight satellites, portable ground stations and real-time secure key exchange” [15]. They brought a complete space-to-ground QKD payload down to 22.7 kg, against roughly 250 kg for Micius, and that figure includes a 200 mm telescope and the pointing system. QubitCore shows that the quantum part need not dominate the mass budget at all (Figure 2).




Figure 2: Payload mass from Micius to QubitCore. Data correspond to Micius, launched in 2016, and the Jinan-1 microsatellite, launched in 2022 [15]. The 22.7 kg of Jinan-1 includes a 14.4 kg pointing system; its QKD terminal alone weighs 8.3 kg [15]. The dark bar is QubitCore, launched in 2026. It holds the photon source and both receivers, but no telescope or pointing system, so the comparison is not like-for-like. The figure shows how small the quantum part of a payload can be.

2.5 The bottleneck and the market behind it

The EU is folding EuroQCI into IRIS², a sovereign constellation of roughly 290 satellites budgeted at EUR 10.6 billion and now expected to enter service around 2030 to 2031 [20, 21]. Eagle-1 (SES and ESA, over EUR 130 million) and SAGA (Thales Alenia, EUR 50 million Phase 1) are the European pathfinders, with national programmes running in Canada, India, Germany, the UK, Spain and Russia, and commercial entrants including SpeQtral, Honeywell, Craft Prospect and IonQ.

Two programmes are cautionary. Arqit raised a USD 400 million deal to list through a shell company on the strength of a satellite QKD plan, kept roughly USD 96 million of it after investors withdrew [23], and then cancelled the space segment entirely. The UK cancelled its ROKS CubeSat before flight. Capital alone does not put a working quantum payload in orbit.

Every one of these programmes needs the same four things: a space-compatible entangled photon source, space-compatible single-photon detectors, a radiation-tolerant time-tagging and processing unit, and the optical and thermal engineering to hold them aligned through launch and thermal cycling. Today each programme builds most of that from scratch, and that is the gap QubitCore fills.

3. QubitCore

3.1 Why the payload runs the full protocol on board

A space quantum communication mission has to solve two hard problems that have little in common. One is quantum optics that survives launch and then works in vacuum, under radiation, across thermal cycles. The other is an optical terminal that finds and holds a link to a ground station while the satellite passes overhead at 7.6 km/s, holding the aim to within microradians. An operational link needs both, but flown together on a first mission they also fail together, and a null result does not say which half was at fault.

QubitCore qualifies the first half on its own. Both photons of each pair go to receivers inside the payload, so the flight exercises the source, the detectors, the time tagging and the BBM92 key protocol with no telescope, no pointing system and no ground segment. Normal count rates, coincidences and key output in orbit therefore qualify the quantum core against the orbital environment, but they say nothing about the space-to-ground channel. This flight does not test its loss budget, pointing accuracy, background light rejection or polarisation alignment, and those are the more difficult part of an operational mission.

Aiming a telescope at a moving ground station and holding it there is a separate discipline with its own suppliers, known as acquisition, pointing and tracking. Its requirements come from the orbit, the aperture and the ground segment rather than from the quantum source. Keeping the two independent lets a customer pair our core with a terminal they already have, or buy one separately, and qualify each on its own schedule.

Parameter	Condition	Typical
Quantum key generation	256-bit keys	15 keys/s with capability up to 500 keys/s
Average power consumption		8 W
Wavelength	Customisable	Photon A 780 ± 5 nm, Photon B 842 ± 5 nm (the two photons of a pair)
Protocol		Entanglement-based BBM92
Working temperature		5 °C to 40 °C
Size and mass		1U (<10 cm x 10 cm x 10 cm cube), 950 g

3.2 Architecture

An entanglement-based QKD link has three roles: Charlie makes the entangled pairs, and Alice and Bob measure them. QubitCore carries all three. Its protocol, BBM92 [7], is named after its three authors and builds the key from the correlations between the two measurements, an idea that goes back to Ekert [8].

A pump laser shines continuously into a nonlinear crystal, which now and then splits one pump photon into two entangled photons at 780 nm and 842 nm. This process is called spontaneous parametric down-conversion. Alice and Bob are two receiving modules, each with its own optics and four silicon single-photon detectors, eight in total. Passive measurement needs four per station: a beam splitter picks one of the two measurement settings at random, and a polarising beam splitter in each arm resolves the two outcomes. Because the randomness comes from the optics themselves, the optical path has no modulators, no high-voltage drivers and no active switching, which for a space payload means fewer components, less power and fewer failure modes.

The two rates in the table come from different limits. The source can make more than it does: at full pump power it produces enough pairs for 500 keys per second. What sets the operating point is the silicon detectors, which stay linear up to a few million counts per second, so we tune the pump to keep them in that range.

Alice shares a module with Charlie and couples to it through free space, while Bob is a separate module coupled by fibre. The payload processing unit is built on an FPGA, a chip programmed to do one job and do it fast. It does the time tagging, correlates events across the eight channels, and runs the comparison and key-generation steps.

The wavelength choice has two consequences. Because the two photons of a pair have different wavelengths, a filter can separate them cleanly instead of losing half of them at a

beam splitter, which matters once the two arms go to two different telescopes. Silicon detectors also work well in this band at temperatures a passive design can hold, so the payload needs no cooler and saves the mass and power one would cost. The drawback is that these wavelengths are wrong for long-haul fibre, which needs 1,550 nm, and that is why we also offer customised wavelengths and fibre or free-space output.

Operation needs no special commands. On power-up the FPGA boots from flash in about 30 seconds, and ten seconds later the pump laser starts and logging begins. Nothing else is required in normal operation, which is why the payload could fly on a standard hosted-payload contract.




Figure 3: First key generation in orbit. Preliminary results confirming that the full chain works in flight: entangled pairs generated, both stations measuring, coincidences found, and BBM92 key production. Five runs between April and August 2026, each of 60 to 300 seconds. QubitCore flies as a hosted payload with limited experiment time, so we took each run under different conditions, including different payload temperatures. The spread between curves reflects those conditions, not a performance limit.

3.3 What flew, and what it qualifies

QubitCore launched on 30 March 2026 aboard SpaceX Transporter-16, riding as a hosted payload on D-Orbit's ION carrier SCV Astounding Alexandra, into a sun-synchronous orbit at about 510 km [17, 18, 19]. A sun-synchronous orbit crosses each point on the ground at the same local time every day. We ran health checks of 60 to 300 seconds, each recording count rates on all eight channels, internal module temperatures, key and bit

generation parameters, coincidence rates between channel pairs, and raw key samples. Carrier telemetry gave us the electrical and thermal context, with current monitored separately for the processing and laser chain and for the detector chain.

Qualified in orbit:

- a miniaturised entangled photon source survives launch and works in the orbital environment
- eight silicon single-photon detectors work on both free-space and fibre-coupled paths
- the FPGA time tagger correlates events and finds coincidences in flight
- the chain from pair generation to key output runs end to end without ground control
- the payload holds its electrical and thermal budget on a commercial hosted-payload interface

Qualifying the quantum payload separately from the space-to-ground link gives it flight heritage of its own, so a customer can pair it with an optical terminal, or buy it for a terminal they already operate, with no re-qualification.

4. Differentiation

4.1 A complete entanglement link in 1U

Other entanglement payloads fly a source and put the receiver on the ground, or fly a source with no protocol at all. SpooQy-1, the closest predecessor, passed a Bell test in orbit, and its authors describe it as a source to combine with optical link technology in later missions [13]. To our knowledge, no other payload has run the complete entanglement-based QKD protocol autonomously in orbit, and that is what makes QubitCore self-verifying: it establishes its own health in flight, with no ground segment at all.

Programme	Payload scale	Reported cost	Status
Qubitrium / QubitCore	1U, 950 g quantum core, no optical terminal	EUR <1 M payload programme, EUR 3.75 M total funding, grants and equity combined	In orbit, Mar 2026
Micius	Dedicated satellite, 631 kg, ~250 kg payload	~USD 100 M	Flown 2016
Jinan-1 microsatellite	22.7 kg payload, telescope included	USD 2 to 5 M (est.)	Flown 2022
SpooQy-1	3U	Academic payload	Flown 2019
OPS-SAT VOLT	16U	EUR 19 M	Target late 2026
SpeQtral-1	16U	USD 15 to 25 M	Target 2027
QEYSSat	Dedicated satellite	CAD 80.9 M	Target late 2026
Eagle-1	Dedicated satellite	> EUR 130 M	Target late 2027

This is not a like-for-like comparison: the larger figures include ground segments, spacecraft buses, operations and, for Eagle-1, an operational service rather than a payload. The direction and the magnitude nonetheless hold. Counting every euro of grants and equity we have ever received, not just the payload budget, we reached orbit on between one thirtieth and one sixtieth of the capital behind the large institutional programmes, and we reached it before most of them will fly. Our own landscape analysis, available on request, covers thirty-seven programmes: most target first flight between 2026 and 2028, several remain conceptual, and two never flew.

4.2 A product, not a mission

QubitCore sits in a catalogue we already sell: entangled photon sources, single-photon detectors in both silicon and InGaAs (silicon for visible and near-infrared light, InGaAs for telecom wavelengths), receiving modules, laser current drivers, picosecond pulsed lasers, high-voltage amplifiers and the QubTag time tagger. Each comes in a space and a terrestrial version, alongside rack-mount QKD systems for ground use [24]. This is a components business in a market where every buyer currently builds bespoke.

4.3 Heritage

Dr Kadir Durak founded Qubitrium in 2020 as a spin-off from Özyeğin University in Istanbul, where he has taught since 2017. He took his doctorate at the National University of Singapore, where he led the team behind the SpooQy-1 nanosatellite entanglement

payload through its critical design review and co-authored the flight demonstration paper [13]. SpooQy-1 is the direct scientific predecessor of this payload class. We now work from Istanbul and Delft, inside the Dutch Quantum Delta ecosystem, and with TNO we are building EPIQ, a second self-contained payload built around a TNO photonic chip that makes entangled pairs at telecom wavelength, targeting late 2026 (Section 5.2).

5. Roadmap: From Qualified Hardware to the Quantum Internet

5.1 The network Qubitrium is building toward

We are building toward a network in four layers, shown in the figure below.




Figure 4: The network Qubitrium is building toward

A **space layer** of CubeSats in low Earth orbit carries entangled sources and distributes entanglement worldwide, over distances and to places fibre cannot reach. A **layer of drones and high-altitude platforms, meaning long-endurance aircraft and balloons**, provides links and timing where radio and satellite navigation are jammed, denied or absent, and covers the gap when cloud blocks an optical downlink. A **ground bridge** of optical stations receives the photons and hands them to the fibre segment. A **fibre network** of QubitEdge nodes around a QubitHub then carries entanglement across a city over partner telecom fibre, which is where the end users sit.

Two things about this architecture matter. The first is that the layers must hand photons to each other, not measurement results. Suppose a ground station measured its photon and the hub then started a fresh pair for the fibre leg: the ground station would hold the key, and the network would be a trusted-relay network again. That is why our first satellite with a downlink sends its photons at 1,550 nm, the wavelength the metro fibre already carries, so the station can pass the photon straight on and QubitHub can route it without ever measuring it.

The second is timing. Entanglement cannot be stored without quantum memory, so a satellite delivers it only while both stations are in view, which for a low orbit is minutes at a time. Keys are different: once made they are ordinary data, and a station can buffer them and spend them over the following day. That is why key services work on a satellite today while the other services in Section 2.1 do not. More satellites shorten the gaps, and quantum memory removes them altogether.

The space layer is the only one that reaches beyond a few hundred kilometres without a trusted relay, and that is why we qualified it first.

5.2 The steps from here

A satellite that makes keys on board produces keys and nothing else. An entangled source, once both photons leave the spacecraft, produces entanglement, which supports everything in Section 2.1. Four steps take us there.

Step	What changes on the payload	What it enables
Now: QubitCore, in orbit	A complete entanglement system in 1U. No optical terminal.	Qualifies the quantum core against the orbital environment.
Next: EPIQ, targeting late 2026	A telecom-wavelength source (TNO's HemiQ photonic chip) and InGaAs detectors. Still no downlink.	Qualifies both halves the fibre layer needs: a source and a detector at 1,550 nm.
Then: optical terminal, one downlink	The source splits each pair into 1,550 nm and 548 nm, both from the same 405 nm pump the present payload uses. The 1,550 nm photon goes down, the 548 nm photon is measured on board. Needs a downlink telescope with serious pointing accuracy.	Entanglement distribution to one ground station, and the photon passes straight into metro fibre.
Then: two downlinks	A second transmitter, because two stations at opposite edges of the footprint are seen from orbit in very different directions [11]. The quantum payload itself does not change.	Entanglement distribution to two ground stations up to about 2,100 km apart.
Then: quantum memory	Memory on board, under development at Qubitrium.	Links beyond a single pass and beyond the footprint limit.

The source is not where the risk sits. Qubitrium already sells an entangled photon source at 1,550 nm and 548 nm as a commercial product with revenue, and QubitCore has shown that our sources survive launch and work in orbit. The real engineering risk is the downlink, and above all, the aiming. The satellite crosses the sky at 7.6 km/s and has to hold its beam on a ground telescope several hundred kilometres away to within a few microradians. That is like threading a needle from the far end of a football field, and holding it there while running.

The photon budget is the other half of the downlink, and the numbers are well-studied. A CubeSat-sized terminal spreads the beam more, which costs roughly a factor of ten at the same range. Better detectors, weaker sunlight at 1,550 nm and larger ground telescopes recover part of that. Still, an operational version needs a well-engineered terminal with microradians pointing accuracy.

The two-downlink step is where the product really changes. Over continental distances, nothing else on the market can put entanglement into two places on the ground at once. Fibre can't reach that far, and a satellite that makes keys on board can't do it at any price.

The field plans its development on a six-stage model, numbered zero to five [16]. Stage zero is a network of trusted relays, where deployed networks including China's sit, and stage one delivers keys end to end without trusting the relay, which is what most announced space QKD programmes target. Stage two distributes entanglement itself, which is what the dual downlink targets. Stage three adds quantum memory, and that is where clock synchronisation, telescope arrays and blind quantum computing become possible.

The market being funded now, by EuroQCI, IRIS², SAGA and their national equivalents, is therefore a stage-one market for keys. A qualified, manufacturable, low-cost entanglement source sells into that market today as a component, and every stage above it needs one. For a buyer or an investor, that is a different risk profile from a pure QKD service, and we think a better one.




Figure 5: The first prototypes of QubitCore on the optical bench.

6. Conclusion

In March 2026 we put a complete entanglement distribution system into low Earth orbit, in a 1U package weighing 950 g and drawing 8 W. Every quantum subsystem worked: the source made pairs, eight detectors counted them, the timing electronics found the coincidences, and the protocol ran end to end and produced keys on its own, with no ground segment involved. We did this on a fraction of the capital behind the programmes we intend to supply, and before most of them will fly.

Space quantum communication missions cost from a few million to around a hundred million each. A large part of the reason is that nobody has been able to buy the quantum

part, so every programme designs, builds and qualifies its own. QubitCore is the first step in changing that, and the components behind it, the sources, detectors, receiving modules and time taggers, are products rather than mission hardware.

From here we add a telecom source, an optical terminal, a second downlink and then memory, building toward a quantum internet whose space layer is small, manufacturable and already flying. The first piece is qualified hardware, in orbit and working.

7. References

- [1] P. W. Shor, SIAM J. Comput. 26(5), 1484 (1997). DOI: [10.1137/S0097539795293172](https://doi.org/10.1137/S0097539795293172)
- [2] C. Gidney, arXiv:2505.15917 (2025). DOI: [10.48550/arXiv.2505.15917](https://doi.org/10.48550/arXiv.2505.15917)
- [3] NIST IR 8547, "Transition to Post-Quantum Cryptography Standards", initial public draft, 12 Nov 2024. Still in draft at the time of writing. <https://csrc.nist.gov/pubs/ir/8547/ipd>
- [4] NIST FIPS 203, 204 and 205, published 13 Aug 2024. <https://csrc.nist.gov/projects/post-quantum-cryptography>
- [5] W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned", Nature 299, 802 (1982). DOI: [10.1038/299802a0](https://doi.org/10.1038/299802a0)
- [6] S. Pirandola, R. Laurenza, C. Ottaviani and L. Banchi, Nat. Commun. 8, 15043 (2017). DOI: [10.1038/ncomms15043](https://doi.org/10.1038/ncomms15043)
- [7] C. H. Bennett, G. Brassard and N. D. Mermin, Phys. Rev. Lett. 68(5), 557 (1992). DOI: [10.1103/PhysRevLett.68.557](https://doi.org/10.1103/PhysRevLett.68.557)
- [8] A. K. Ekert, Phys. Rev. Lett. 67(6), 661 (1991). DOI: [10.1103/PhysRevLett.67.661](https://doi.org/10.1103/PhysRevLett.67.661)
- [9] S.-K. Liao et al., "Satellite-to-ground quantum key distribution", Nature 549(7670), 43 (2017). DOI: [10.1038/nature23655](https://doi.org/10.1038/nature23655)
- [10] J. Yin et al., "Satellite-to-Ground Entanglement-Based Quantum Key Distribution", Phys. Rev. Lett. 119(20), 200501 (2017). DOI: [10.1103/PhysRevLett.119.200501](https://doi.org/10.1103/PhysRevLett.119.200501)
- [11] J. Yin et al., "Entanglement-based secure quantum cryptography over 1,120 kilometres", Nature 582(7813), 501 (2020). DOI: [10.1038/s41586-020-2401-y](https://doi.org/10.1038/s41586-020-2401-y)
- [12] Y. Liu et al., Phys. Rev. Lett. 130(21), 210801 (2023). DOI: [10.1103/PhysRevLett.130.210801](https://doi.org/10.1103/PhysRevLett.130.210801)
- [13] A. Villar et al., "Entanglement demonstration on board a nano-satellite", Optica 7(7), 734 (2020). DOI: [10.1364/OPTICA.387306](https://doi.org/10.1364/OPTICA.387306)

- [14] Y.-A. Chen et al., Nature 589(7841), 214 (2021). DOI: [10.1038/s41586-020-03093-8](https://doi.org/10.1038/s41586-020-03093-8)
- [15] Y. Li et al., "Microsatellite-based real-time quantum key distribution", Nature (2025). DOI: [10.1038/s41586-025-08739-z](https://doi.org/10.1038/s41586-025-08739-z)
- [16] S. Wehner, D. Elkouss and R. Hanson, Science 362(6412), eaam9288 (2018). DOI: [10.1126/science.aam9288](https://doi.org/10.1126/science.aam9288)
- [17] D-Orbit, "Wayfinder" mission page, March 2026. <https://dorbit.space/wayfinder-march-2026/>
- [18] SpaceNews, "SpaceX launches Transporter-16 rideshare mission", 30 March 2026. <https://spacenews.com/spacex-launches-transporter-16-rideshare-mission/>
- [19] The Quantum Insider, "A Quantum Payload Reaches Orbit: Commercial Quantum Communication Is on the Horizon", 14 April 2026. <https://thequantuminsider.com/2026/04/14/a-quantum-payload-reaches-orbit-commercial-quantum-communication-is-on-the-horizon/>
- [20] European Space Agency, "European Quantum Communication Infrastructure (EuroQCI)", ARTES 4.0 programme page. <https://connectivity.esa.int/artes-4-0-programme-overview/optical-quantum-communications-scyllight/european-quantum-communication-infrastructure-euroqci>
- [21] EUSPA, "OBSERVER: A quantum secure future with IRIS²". <https://eu-space.europa.eu/news/observer-quantum-secure-future-iris2>
- [22] TeleGeography, "Submarine Cable FAQs", citing US Federal Communications Commission figures showing satellites at 0.37% of US international capacity. <https://www2.telegeography.com/submarine-cable-faqs-frequently-asked-questions>
- [23] Arqit Quantum Inc., Form POS AM, FY2021, US Securities and Exchange Commission, 16 December 2021. <https://www.sec.gov/Archives/edgar/data/1859690/000110465921150291/arqq-20211216xf1.htm>
- [24] Qubitrium, company and product information. <https://qubitrium.tech/about> and <https://qubitrium.tech/products> ; ACT Venture Partners, "Why Invested in Qubitrium". <https://actvp.vc/stories/tpost/gj15ypkh71-why-invested-in-qubitrium>
- [25] M. Krelina, U. Tefek, Z. C. Seskir and K. Durak, "Quantum Protocols for Time Synchronisation and Distribution: A Critical Assessment", arXiv:2604.10243 (2026). DOI: [10.48550/arXiv.2604.10243](https://doi.org/10.48550/arXiv.2604.10243)

[26] D. Gottesman, T. Jennewein and S. Croke, "Longer-Baseline Telescopes Using Quantum Repeaters", Phys. Rev. Lett. 109, 070503 (2012). DOI: [10.1103/PhysRevLett.109.070503](https://doi.org/10.1103/PhysRevLett.109.070503)