OPEN LUNAR FOUNDATION

Space Time Card: A Near-Term Bridge and Long-Term Resilience Layer for Lunar Time & Coordination

Philip Linden
Ashley Kosak
2026

Space Time Card: a near-term bridge and long-term resilience layer for lunar time & coordination

Philip Linden
philiplinden@openlunar.org

Ashley Kosak
ashleykosak@openlunar.org

Executive Summary

Large-scale space capabilities rarely emerge fully formed. As constrained, standardized building blocks gain relevance through repetition, interoperability, and declining marginal cost, progress is driven less by individual platform performance and more by the accumulation of flight heritage and the stabilization of interfaces. Once a critical mass is reached, these architectures rapidly transition from experimental curiosities to default infrastructure. This pattern has appeared in multiple domains of space systems engineering— including CubeSats where early emphasis on openness, standardization, and robust hardware architecture proved more durable than early optimization for peak capability and ushered in a paradigm shift in satellite operations.

Timekeeping in lunar operations is especially well-positioned to benefit from such an approach. Building large geospatial positioning system satellites requires relatively large financial investments while the lunar economy is still in its infancy. On-board solutions allow timekeeping to advance in parallel with a growing ecosystem of lunar spacecraft, improve landing system performance, and could be achieved at a significantly lower cost as described in the paper below.

The Epoch Project is a collaborative effort by Philip Linden and Ashley Kosak, with support by the Open Lunar Foundation and Olivia Linden. The project aims to advance the Space Time Card, an independent timing system with aerospace-grade precision and accuracy. The objective is to bridge the gap of external timing sources around the Moon. It reframes lunar timekeeping from a specialized subsystem into shared, enabling infrastructure. This is in line with the Open Lunar Foundation's approach of investing in infrastructure which supports collaborative, co-invested solutions which align with sector readiness and present relevance.

By combining well understood physics, flight proven synchronization techniques, accessible hardware, and open interoperability principles, Epoch provides a practical path for missions to achieve autonomy today while contributing to a resilient lunar positioning, navigation, and timing (PNT) ecosystem tomorrow. The result is not a single authoritative clock or constellation, but a network of independently validated nodes whose collective behavior strengthens with each deployment. As lunar activity scales, this incremental, flight-proven approach offers a durable foundation for coordination, navigation, and reference time beyond Earth orbit.

This paper delivers the building block: a standalone, bootstrapped timing node. Diverse and decentralized networks naturally emerge, grow, and adapt from large populations of nodes.

The Space Time Card is designed to operate within contemporary lunar interoperability efforts such as the NASA LunaNet and ESA's Moonlight frameworks and related international navigation initiatives. Rather than competing with future lunar navigation services, The Space Time Card provides a practical near term bridge. It enables missions to operate autonomously today while remaining compatible with shared infrastructure as it emerges, and it retains long term value as a resilience and holdover layer even once external services mature.

The near-term performance objective is to maintain sub-50 nanosecond time deviation over 48 hour autonomous holdover periods without external synchronization. The immediate benefit comes from carrying multiple onboard oscillators, cross-disciplined against one another and combined by consensus such that the system is maintaining stable and high precision time internally with minimal drift. The distributed network becomes an additive capability over time, the prerequisite is establishing the ubiquity of on-board precision time.

OPEN LUNAR
FOUNDATION

Table of Contents

1. An Open, Incremental Approach to Lunar Time

1.1 Why Lunar Timekeeping Must Move Onboard

Lunar exploration is entering a period of rapid expansion. Government programs and commercial landers and orbiters are launching at a higher cadence than in the previous decade. Unlike Earth orbit, where GNSS (Global Navigation Satellite System) provides continuous timing and navigation services, lunar missions must account for intermittent communications, light-time delays of 1.3 seconds or more, and limited shared infrastructure during operations. As mission count increases, reliance on centralized ground-based tracking faces scaling issues because ground-based tracking resources are necessarily finite and scheduled across missions. The approximately 2.6 second round-trip light-time between Earth and Moon constrains high-bandwidth closed-loop control from Earth. As mission cadence increases and operations diversify, spacecraft must assume greater responsibility for time-sensitive estimation and coordination onboard.

Precision timekeeping underpins navigation, coordination, and measurement in deep space. A trusted time base allows received signals to be interpreted onboard, measurements to be correlated across platforms, and state estimation to run autonomously. In cislunar space, the current absence of GNSS prevents missions from onboard navigation autonomy and requires building bespoke timing subsystems that are costly, difficult to validate, and not reusable across programs (Manning, 2023). As a result, missions incur higher integration and validation costs, face reduced reuse across programs, and operate with a lower autonomy ceiling than would otherwise be achievable.

The Epoch Project is exploring a different path. By making onboard precision timekeeping accessible, interoperable, and repeatable, lunar positioning, navigation, and timing (PNT) capability can grow organically and immediately. Rather than deploy a singular infrastructure, The Epoch Project proposes deploying an infrastructure in successive nodes, so that each deployment lowers integration risk for subsequent missions as flight heritage accumulates and characterization of performance envelopes become empirical. The output being a robust lunar timing network which emerges incrementally rather than waiting for centralized infrastructure to be deployed.

This paper describes the technical and governance foundations of Epoch, presents the Space Time Card reference implementation, and outlines validation approaches and deployment scenarios. Our goal is to complement the future lunar navigation efforts such as NASA's LunaNet and ESA's Moonlight, which are primarily focused on communications relays, interoperability standards, and the transport of timing and navigation data (IOAG and ICG 2025).

We intend to deliver a near-term bridge and long term resilience layer by enabling locally available time and coordination capabilities which are useful before, during, and after external lunar infrastructure matures.

1.2.1 Relativistic Time in Cislunar Space

Time does not advance at the same rate on the Moon as it does on Earth or anywhere in between. For most terrestrial applications, this is a negligible effect. In extremely sensitive applications like high-frequency trading and GNSS satellites orbiting the Earth, it is critical to correct for this phenomenon called time dilation. Differences in gravitational potential and orbital velocity cause a clock at the lunar surface to advance faster than an Earth-based clock by approximately 56 microseconds per day. Left uncorrected, this rate offset accumulates into tens of kilometers of equivalent range error over mission-relevant timescales.

The dominant effect is characterized by a dimensionless rate factor of approximately 6.5×10^{-10}, with additional periodic variations driven by orbital eccentricity and geometry. While small on human timescales, these effects are significant for precision navigation and must be modeled explicitly rather than treated as secondary corrections (Ashby, 2024).

1.2.2 Capabilities and Limits of Compact Precision Clocks

Recent advances in chip-scale atomic clocks (CSACs) have changed what is feasible for small missions. Devices such as the Microchip SA.65 provide strong short-term stability in compact, low-power packages. The SA.65 occupies less than 17 cm^3 and consumes under 120 mW to deliver Allan deviation of approximately 3×10^{-10} at 1 second (Microchip Technologies). These capabilities expand participation in deep space timing beyond flagship missions. Early missions incorporating CSAC-class timing, such as the CAPSTONE mission, incurred substantial non-recurring engineering costs; however, this gap is expected to narrow as these integration efforts become more standardized and flight heritage increases.

However, CSACs alone cannot maintain nanosecond-level accuracy indefinitely. Frequency stability degrades with aging and environmental sensitivities. Without periodic resynchronization, a CSAC-based timing system can drift by hundreds of microseconds within days. This limitation motivates an ensemble approach.[1] By combining multiple oscillators and applying statistical weighting and fault detection algorithms, stability and resilience beyond that of any single oscillator can be achieved (Byagowi et al., 2022).

[1] Throughout this paper, the term *ensemble* refers to a statistical combination of multiple oscillators, not machine learning ensemble methods.

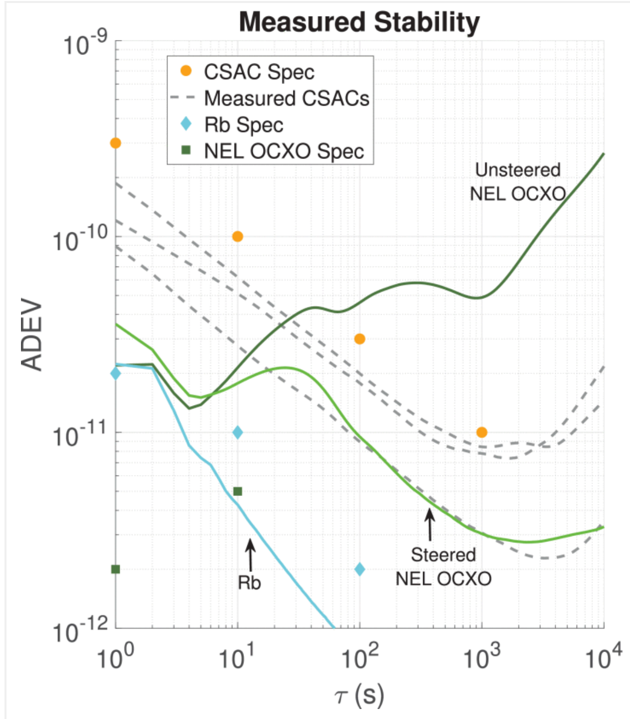

Figure 1. Benchtesting results for ADEV of Steered OCXO to implicit ensemble mean (IEM) of three chip scale atomic clocks. Allan deviation(Y-axis) as a function of averaging time τ (X-axis) for several representative oscillator classes and configurations (Flood, 2023).

Ensemble techniques also enable graceful degradation: if one oscillator fails or shows anomalous behavior, the system can detect and exclude it without losing time continuity. Heterogeneous ensemble performance has been experimentally demonstrated by Flood et al. at the Air Force Research Laboratory in collaboration with the University of Colorado Boulder (Flood et al., 2023). In their testbed, CSACs and OCXOs were combined and referenced to a rubidium frequency standard, with frequency control implemented using a low-noise 18-bit DAC. The results shown in Figure 1 demonstrate stabilized Allan deviation performance relative to single-oscillator configurations.

1.3 Interoperability as Foundational Infrastructure

Lunar timekeeping is also a governance and ecosystem challenge. No single operator should control lunar time, and no single vendor or architecture should gate participation. Centralized stewardship of reference time introduces structural concentration. A single authoritative source creates a single point of technical, operational, or policy failure. Service outages, prioritization decisions, access controls, or interface changes propagate system-wide effects. In a multi-actor lunar environment involving government, commercial, and international participants, such concentration may also create dependency on a single steward for mission-critical coordination, potentially being the decision maker for access or limiting architectural diversity.

An open core approach separates control of time from control of services by defining a shared, interoperable baseline for lunar timing and coordination, which allows multiple actors to build, operate, and fund differentiated capabilities on top of that foundation. Interoperability enables cooperation across agencies, companies, and nations without requiring shared ownership or centralized control.

Open and neutral standards align incentives across stakeholders:

- Academic teams gain access to flight-relevant problems and contribute validation data
- Industry benefits from reduced integration cost and shared evidence of performance
- International partners can participate without ceding operational independence
- Small missions can adopt proven designs without starting from scratch

These dynamics mirror prior successes in precision timing systems, where open reference designs accelerated adoption. The Open Compute Project's Time Appliances working group has driven down cost and integration barriers for datacenter timing through shared schematics, firmware, and test procedures.(Open Compute Project) CERN's White Rabbit project established sub-nanosecond synchronization across distributed systems by publishing complete reference implementations and fostering a vendor-neutral collaboration. (CERN, 2024)

Epoch complements emerging lunar interoperability efforts such as NASA's LunaNet framework for supplemental navigation services in PNT and ESA's Moonlight initiative (NASA 2025). Rather than competing with planned lunar navigation services, Epoch provides:

1. **A near-term bridge:** Missions launching in this decade benefit from onboard timing and coordination capabilities which reduce reliance on continuous ground station support and enable autonomy, independent of operational LunaNet or Moonlight relay services.

2. **A resilience layer:** Even after external infrastructure is deployed, onboard timing provides autonomy during outages or when links are unavailable

3. **A validation platform:** Flight data from Epoch nodes can inform requirements and architectures for future centralized services

To promote interoperability, the Epoch Project is aligning its operating standards with the relevant elements of the LunaNet Operability Specification (LNIS), including message formats for two-way time transfer (TWTT) observables. As infrastructure emerges and matures, protocol evolution will likely occur through established standards bodies and community forums such the Consultative Committee for Space Data Systems and other consultative processes. We intend to operate with these interoperability frameworks rather than inventing a new standard. Governance questions specific to lunar time scale definition and multi-actor coordination are being actively explored through forums such as CCSDS and the Inter-Agency Working Group on Cislunar PNT (ICG-IOAG).

1.3.1 Buy, Blend, or Brew It Yourself

The open-source community defines "free as in free beer" as in *it is free to drink, but you don't get to control how it's made*. If we shift the perspective from the *consumer* to the *producer*, the analogy has a new meaning:

> *Beer costs money to buy and to make, but there is a spectrum of investment, time and effort in how you get your own drink.*

Even when interfaces and designs are openly available, system builders face a range of choices in how they adopt them. Some purchase ready made solutions, others combine existing components, and some build entirely from first principles. These choices represent tradeoffs between upfront cost, development time, and control rather than differences in access or permission.

Choose a bottle from a wide selection of brands and flavors and drink it right away. Start from scratch with raw ingredients, for complete control over the flavor or because it's more economical to buy ingredients in bulk, then brew it yourself for a few months. Split the difference and buy multiple flavors, then blend them together at home to get your own perfect flavor.

Buy it ready-made. Manufacturers anticipate what customers need. Standard options are cheap; boutique solutions get expensive. Off-the-shelf always either overkill with capabilities you don't need or underkill with other systems handling the missing features. But when your time costs more than the premium, buying is the right call. A mission with a tight schedule and proven requirements can integrate a complete Space Time Card module and move on. A mission with a very specific use case and flexible schedule can design the perfect module in-house or subcontract something custom.

Blend existing options. Achieve compositional control by combining, substituting, and evolving components within a shared structure without fragmenting the interfaces. Swap a different oscillator onto the daughter board. Run modified firmware on the same hardware. Combine an off-the-shelf CSAC with a custom stability control system. The goal is consistency rather than peak performance, and the ability to spin up new configurations with minimal rework. Blending allows the design specific requirements and constraints *as a system* by designing how the pieces work together instead of making the all-in-one solution from scratch.

Brew it yourself. DIY is more than repeating the typical approach. Start with the reference design as a kit. Follow the same architecture with cheaper or more available parts. Use the same parts with a different architecture. Make incremental changes to suit your requirements. Or fork the repo and experiment from scratch. Each approach trades dollars for effort at a different rate. Brewing makes sense when nothing on the market fits, when the learning itself is the goal, or when you need capabilities that don't exist yet or could be added later.

1.3.2 Ecosystem Analogs and the Case for Open Reference Design

In terrestrial precision timing, the Open Compute Project Time Appliances Project and its Time Card and Open Time Server efforts illustrate the pattern. The project frames the Time Card as a traceable time of day source that can serve both directly attached systems and distributed systems, and it explicitly emphasizes standardized architecture, interfaces, and figures of merit to characterize performance (Open Compute Project, 2021). The associated open implementation work, published as an open repository, positions Time Card as an open hardware and software approach that can turn commodity servers into time servers with GNSS input and holdover oscillator support, creating a reproducible reference that others can build on rather than reverse engineer (Byagowi, 2021). This is valuable because the community aligns on common interfaces, common measurements, and shared validation paths.

White Rabbit provides the complementary precedent for interoperability in distributed system synchronization. CERN describes White Rabbit as an Ethernet based timing technology that supports precision time tagging and triggering in large installations, and the Open Hardware Repository presents White Rabbit as an open hardware and software ecosystem that achieves sub nanosecond accuracy and picoseconds level precision in distributed systems. Beyond the existence of an open implementation, White Rabbit demonstrates the governance and compatibility layer that makes interoperability real: well defined profiles, transparent performance expectations, and community maintained evolution, including the formal White Rabbit Collaboration and alignment with the IEEE 1588 2019 High Accuracy Profile (CERN, accessed 2026).

The Epoch Project can mirror the conservative parts of these models. The compatibility promise should be framed around stable interfaces and conformance tests, not around mandating a single hardware choice. Publishing timing interfaces, measurement definitions, and flight relevant conformance tests enable different vendors and different clock grades to interoperate at defined performance tiers. Integrators can verify behavior in the lab before flight. This approach translates directly into faster iteration because teams can depend on validated interfaces and harnesses.

1.4 Case Studies & Lessons from the Past

Flight heritage sets the stage for ongoing development of infrastructure within the aerospace industry. As the Space Time Card continues to develop, it's important to recall the missions and platforms which enable low-cost precision timekeeping infrastructure. We must also consider how these missions demonstrate that meaningful operational autonomy in cislunar space is both achievable and valuable.

1.4.1 CAPSTONE: Small Platform Flight Heritage for CSAC-class Timing

The CAPSTONE mission provides the clearest flight precedent for onboard precision time enabling autonomy in cislunar space. Advanced Space launched its Cislunar Autonomous Positioning System, CAPS, on June 28, 2022 as part of the CAPSTONE mission, a 12U CubeSat-class spacecraft with a mass of approximately 25 kg. CAPS successfully demonstrated peer-to-peer navigation techniques, enabling relative tracking between spacecraft and advancing autonomous navigation capabilities beyond continuous ground based support. By reducing dependence on Earth based tracking, CAPS showed that small spacecraft can perform meaningful onboard navigation in cislunar space (Agasid 2024).

The CAPSTONE Mission serves as a permission slip for small platforms operating in lunar-relevant orbits. Flown as a compact spacecraft in a near rectilinear halo orbit, CAPSTONE demonstrated that autonomy is achievable within the size, power, and integration constraints typical of small commercial missions.

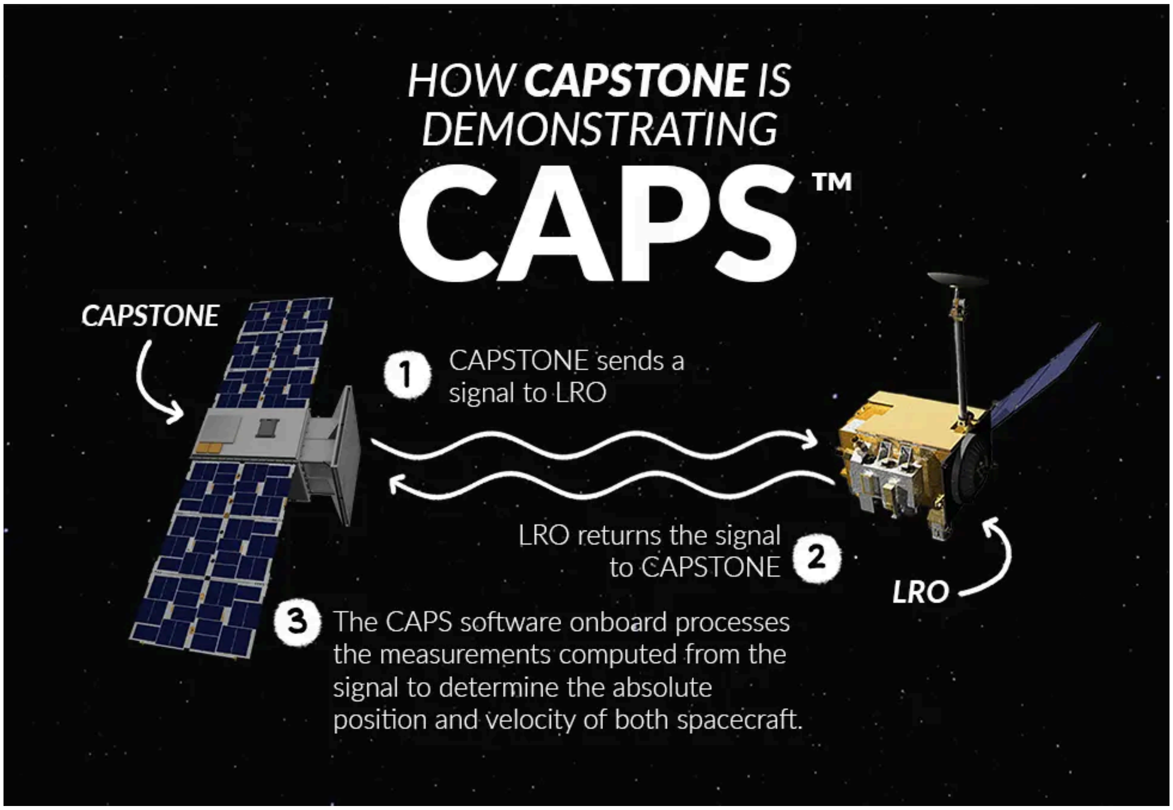

Figure 2. CAPS Demonstration high level overview. (Advanced Space, 2022)

At a system level, CAPSTONE combined three tightly coupled elements on a constrained platform:

1. An onboard CSAC-class timing source.
2. A spacecraft radio and communications system.
3. Onboard navigation and estimation software.

Together, these elements enabled exploration of one way radiometric navigation concepts and reduced exclusive reliance on two way Earth based tracking.

The lesson from CAPSTONE is not that integration was trivial, but that it was achievable within constraints typical of university and small commercial missions. When a spacecraft can trust its onboard time base, received observables become actionable inputs rather than raw data awaiting ground correlation. This enables navigation services that do not require continuous ground contact and directly supports the claim that interoperable onboard precision timekeeping for lunar missions is achievable now without bespoke timing subsystems or constant two-way ground dependence.

Achievement of a successful CAPSTONE mission with a constrained platform also implies accessible scalability. As a 12U CubeSat, CAPSTONE sits in a similar size and complexity class as many small commercial platforms. The architecture is achievable under tight constraints. When this architecture is documented and repeatable, each integration effort becomes reusable infrastructure rather than a one off reinvention. CAPSTONE lays the groundwork for scalability of onboard precision timekeeping by small teams. For the Space Time Card, CAPSTONE demonstrates that CSAC-class timing can be integrated on constrained small platforms and can materially support onboard autonomy, establishing a practical precedent for repeatable, small-mission precision timekeeping.

1.4.2 Deep Space Atomic Clock: A Benchmark for Onboard Time and the Autonomy Ceiling

NASA's Deep Space Atomic Clock (DSAC) is the clearest modern benchmark for what "excellent onboard time" can look like in flight, and why it changes navigation operations. DSAC's purpose was to push reliable timing stability into a flight compatible instrument, so future spacecraft can navigate with less frequent dependence on Earth based time transfer and continuous scheduling through ground networks (NASA 2023).

DSAC flew as a NASA technology demonstration, hosted on General Atomics' Orbital Test Bed spacecraft and launched on June 25, 2019. The payload used a mercury-ion trap atomic clock, a fundamentally different technology than the rubidium physics packages in CSACs. The performance is exceptional, but the development cost (>$1M), mass (19 kg), power (56 W), and integration complexity place it beyond reach for most lunar missions.

Published results report long term fractional frequency stability of about 3×10^{-15} at 23 days with no drift removal, and a measured linear drift of about 3×10^{-16} per day, exceeding prior space clock stability by up to roughly an order of magnitude. NASA also describes the on orbit record operationally, reporting a time deviation under four nanoseconds after more than 20 days of operation.

The DSAC results explicitly connect improved clock stability to enabling one way navigation, where signal delay times can be measured in situ rather than requiring a two way turnaround for each measurement cycle.(Stuart, 2020) When the spacecraft can maintain a sufficiently stable onboard time base, it can treat received radiometric observables as actionable inputs for state estimation without coupling every routine measurement to a complete two way ground exchange.

DSAC defines the upper bound of what exceptional onboard time can enable: stable onboard time that unlocks high quality one way observables and reduces routine dependence on constant ground cadence, while staying accessible to small teams. DSAC shows what outstanding onboard time can enable.

1.4.3 Global Positioning System: The Hero and the Antihero of PNT

Modern society depends on GPS for both positioning and precision timing, underpinning systems from financial markets to logistics and telecommunications. Conceived in the early 1970s primarily to meet United States military navigation and timing needs, GPS introduced a model in which authoritative time is generated and controlled by a centrally managed satellite constellation and distributed to passive users. Over time, policy decisions expanded civilian access to this capability, notably after the 1983 downing of Korean Air Lines Flight 007, which prompted U.S. policy shifts toward eventual civilian access to GPS signals and highlighted the risks of restricted navigation infrastructure. While this evolution democratized access to space based timing, it did not change the underlying architecture. Authority over time remained centralized, sovereign, and dependent on continuous space and ground segment stewardship.

The Space Time Card inverts this architectural assumption. Rather than relying on a distant authoritative system to broadcast time, onboard precision timekeeping places high quality time generation and maintenance directly on the spacecraft. In the lunar environment, where latency, intermittent visibility, and sovereign dependence complicate centralized navigation services, this shift enables timekeeping to function as an onboard capability rather than an externally granted service. The result is not merely broader access, but a redistribution of timing authority itself.

Centralized reference time systems also introduce systemic risk when used outside their original design envelope. Dependence on a single authoritative source can propagate errors, degrade autonomy during outages or disruptions, and constrain mission operations when ground contact is unavailable.

> *"The GPS system cannot serve as a sole source for position location or precision timing for certain critical applications. Public policy must ensure that safety is maintained, even in the event of the loss of GPS. [...]*
>
> *Backups for positioning and precision timing are necessary for all GPS applications involving the potential for life-threatening situations or major economic or environmental impact." (Volpe, 2001)*

By contrast, distributed onboard clocks with cross validation reduce reliance on any single reference and reduce reliance on routine two way ground ranging, allowing missions to maintain timing integrity during extended periods without continuous ground synchronization. As lunar activity scales, such systems can further support peer to peer time confirmation across multiple spacecraft, forming a resilient network in which timing accuracy emerges from redundancy and cooperation rather than from a single controlling node.

From a funding and programmatic perspective, independent onboard timing offers an immediate, low cost path to reducing dependence on large scale lunar GNSS infrastructure that remains expensive, complex, and schedule uncertain. The initial GPS constellation required on the order of tens of billions of dollars to design, deploy, and sustain, reflecting both its scale and its role as a centralized, sovereign system. Proposed lunar navigation constellations, including CubeSat based approaches, meaningfully reduce cost relative to GPS but still require coordinated development, multi launch deployment, and years of program execution before full capability is realized.

Onboard precision timekeeping decouples lunar PNT capability from these long horizon dependencies. By placing high quality time generation and holdover directly on spacecraft, missions gain access to timing functionality independent of external constellation readiness, launch schedules, or service availability. This enables near term mission autonomy while allowing navigation and timing capability to scale incrementally as additional infrastructure comes online. Rather than positioning onboard timing as a replacement for future lunar GNSS, this approach functions as a robust stepping stone: it reduces early mission risk, lowers the barrier to entry for new lunar actors, and complements eventual constellation based services when they mature

2. Introducing the Space Time Card

This whitepaper presents a concept and roadmap for the *Space Time Card*, a plug-and-play time appliance for spacecraft. It brings interoperability, incremental progress, and access to precision timing beyond low Earth Orbit. The Epoch project's Space Time Card supports innovation and diversity in the technological landscape for PNT in a way that does not compromise interoperability or architectural freedom for the sake of adoption.

The Space Time Card architecture reflects deliberate choices about how technology spreads in practice. Open standards enable broad participation only when they accommodate the spectrum of ways teams actually adopt new systems. A university team with limited budget and extensive schedule operates differently from a commercial integrator with tight deadlines and proven requirements, for example. The Space Time Card separates stable interfaces from implementation choices so these diverse users can participate without fragmenting compatibility or forcing unnecessary convergence.

2.2 System Architecture

The Space Time Card operates as a dedicated network device on the internal ethernet network of the host spacecraft. The host connects to an internal system-on-chip (SoC) over ethernet, where the Space Time Card presents itself as a PTP (Precision Time Protocol) Leader. This reference design in Figure 3 assumes that the host is capable of operating without the Space Time Card and uses ethernet as the network architecture. Derivative Space Time Cards could be designed to be compatible with other network interfaces, such as SpaceWire.

The SoC responds to all commanding, configuration, and telemetry requests from the host's primary computer. The SoC serves as a network device visible to the host. In other words, the host does not have direct access to the card and therefore doesn't have to manage any software functions or operations internal to the device. The SoC abstracts the interface to the host so the host can be agnostic of the Space Time Card internals. Internally, the SoC is integrated with an FPGA that does the heavy lifting with respect to timekeeping. The host does not have direct access to the FPGA or oscillator.

The Space Time Card is an architectural descendant of the Open Compute Project Time Appliances Project (OCP-TAP) Time Card (Byagowi et al. 2022) with a limited feature set, borrowing hardware and software designs from the open-source reference design (OCP Time Appliances Project 2024) and other open-source derivatives (Safran Navigation & Timing 2023).

While the logical AXI-bridge architecture and software-visible time registers defined in the OCP-TAP specification are preserved, we have modified the physical layer for vacuum and high-vibration environments. Specifically, we replaced the standard PCIe interface with a ruggedized Ethernet/PTP stack to serve as the vehicle's "PTP Leader". This ensures that while the form factor is optimized for space, the data models and firmware interfaces remain compatible with the broader OCP ecosystem.

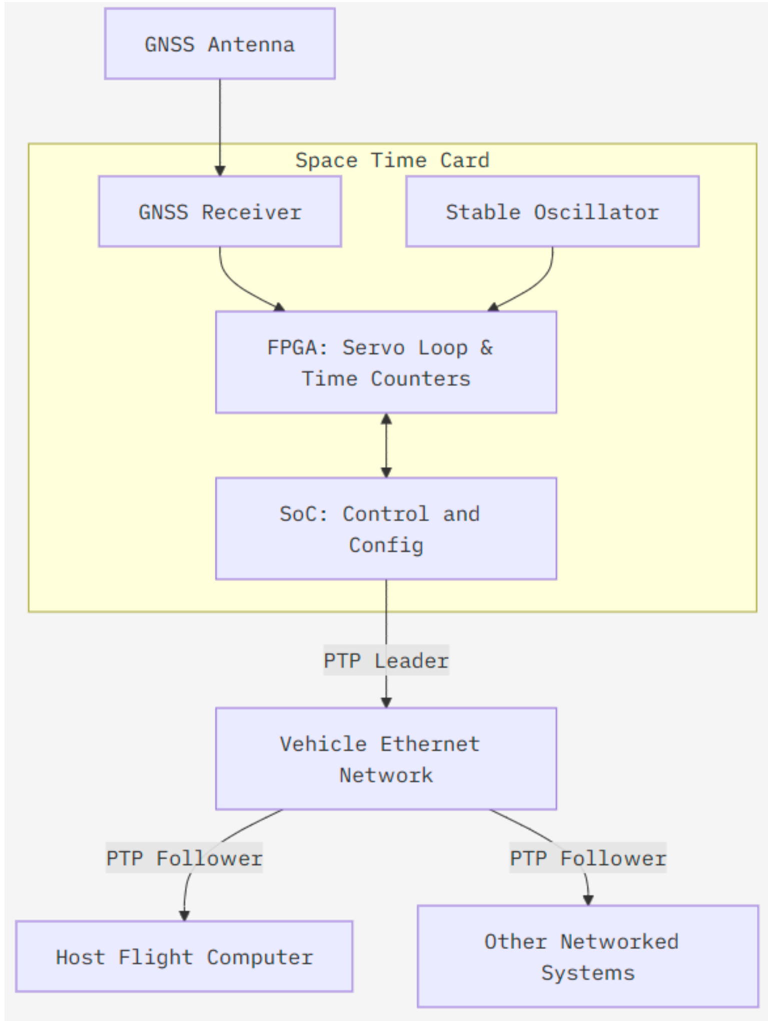

Figure 3. High-level overview of the Space Time Card system Architecture.

The core architecture of a Space Time Card consists of:

- A microcontroller or CPU that interfaces with the host for all data communication.
- Support for 1–4 precision oscillators (CSAC, TCXO, OCXO, MAC, and/or GNSS-disciplined).
- Standardized software-visible time registers (hardware timestamping at interfaces).
- Separation between oscillator hardware and discipline logic.

This architecture enables missions to trade performance, power, and cost without redesigning host interfaces. A university mission might fly a single CSAC, while a commercial navigation mission might integrate a redundant ensemble of atomic and crystal oscillators. Likewise, development efforts and experimentation may be localized to a specific aspect of the architecture without major rework to integrate the Space Time Card with the host spacecraft.

2.2.1 Space Environment

The Space Time Card specifically targets the use case of providing precision timing to a spacecraft in Earth orbit or deep space. Exposure to the space environment is expected to be abated by the spacecraft bus and enclosure somewhat, but the Space Time Card will still be subject to extreme total radiation doses, thermal, and electromagnetic conditions. Managing the local thermal and electromagnetic environment is the responsibility of the host vehicle. The Space Time Card's sensitivity to radiation, thermal management, and other factors shall be characterized during flight qualification in Low Earth Orbit (LEO) initially, then characterized for cislunar conditions in subsequent qualification campaigns.

LEO is the initial target environment for Space Time Card qualification. LEO satellites typically rely on GNSS for positioning, navigation, and timing. The Space Time Card provides resilience to GNSS signal instability or interference by maintaining precision timing through holdover. With frequent small satellite launches and decreasing access-to-orbit costs, LEO offers an opportunity to iterate on the design through successive flight experiments while serving an immediate market need. Development toward cislunar environments continues in parallel.

Commercially available precision oscillators on the market spanning orders of magnitude in cost and size, weight[2] and power (SWaP). Table 2 and Table 3 contextualize a number of space-qualified atomic clocks to illustrate the full performance-SWaP-cost trade space.

Table 1. Environmental characteristics across mission environments.*

	Low Earth Orbit (LEO)	Geostationary Earth Orbit (GEO)	Cislunar Transfer	Low Lunar Orbit	Lunar Surface
Env. Temp. Extremes[3]	−40°C to +85°C	−65°C to +150°C	−150°C to +150°C	−40°C to +85°C	−180°C to +130°C
Eclipse Duration	~35 min/orbit	~72 min/day (equinox seasons)	Variable (up to hours)	30-45 min/orbit	~14 Earth days
Total Ionizing Dose (TID), 5 yr	5–50 krad(Si)	50–100+ krad(Si)	10–50 krad(Si)	2–15 krad(Si)	2–10 krad(Si)

[2] The colloquial term SWaP includes *weight* but space missions consider *mass*. We use "weight" here to stand in for mass in order to preserve the shorthand term.

[3] The desired operating temperature is approximately -10°C to +80°C. Thermal regulation is assumed to be handled by the host spacecraft. When the environmental temperature is lower than the operational minimum, additional power is required for heaters. When it is above the operating temperature, it is difficult to reject excess heat by radiation alone. As environmental temperature diverges from the desired range, thermal management increases in complexity, SWaP, and cost.*LEO*: NASA-HDBK-4002A (2017), Wirthlin et al. (2013).
GEO: ECSS-E-ST-10-04C (2008), Ginet et al. (2013).
Cislunar: Mazur et al. (2011), Guo et al. (2019).
Lunar orbit: Spence et al. (2010), Minow et al. (2019).
Lunar surface: Williams et al. (2017), Paige et al. (2010), Minow et al. (2019).

Single Event Upset (SEU) Rate	10^{-7}–10^{-6}/bit/day	10^{-6}–10^{-5}/bit/day	10^{-7}–10^{-5}/bit/day		
Dominant Radiation Sources	high energy protons & electrons (solar weather)	Outer belt electrons, solar weather	Cosmic rays, solar weather, Van Allen transits	Cosmic rays, solar weather	Cosmic rays, solar weather, albedo neutrons

Table 2. Typical cost of space-rated atomic clocks across a range of precision.

	Cost	Description
Microchip SA.65 (CSAC)	< $10,000	Radiation-tolerant variant
Safran mRO-50 (Rb)	$15,000–$30,000	Military/aerospace ruggedized
Excelitas GPS RAFS (Rb)	$200,000–$500,000	Full GNSS space heritage, extensive qualification
DSAC (Hg^+)	>$1,000,000	NASA flagship tech demo, not optimized for size, mass, or power

Table 3. Typical performance of space-rated atomic clocks across a range of precision.

		Microchip SA.65 (CSAC)	Safran mRO-50 (Rb)	Excelitas RAFS (Rb)	NASA DSAC (Hg^+)
Frequency Stability (ADEV)[4]	τ = 1 s	3×10^{-10}	$\leq 6 \times 10^{-11}$	2×10^{-12}	—
	τ = 10 s	1.3×10^{-10}	$\leq 1.9 \times 10^{-11}$	6×10^{-13}	—
	τ = 100 s	4×10^{-11}	$\leq 6 \times 10^{-12}$	2×10^{-13}	2×10^{-13}
	τ = 24 h	3×10^{-11}	$\leq 1 \times 10^{-11}$	$\leq 5 \times 10^{-14}$	$\leq 2 \times 10^{-15}$
	τ = 552 h (23 days)	2×10^{-10}	5×10^{-11}	3×10^{-11}	3×10^{-15}
SWaP	Power	0.12 W	0.5 W	≤14 W	56 W
	Volume	0.017 L	0.52 L	4.2 L	19 L
	Mass	35 g	80 g	6.4 kg	19 kg
Space Heritage		CAPSTONE (2022)	Military ruggedized (2020s)	GPS IIA/IIR/IIF (1990-2016)	STP-2 Orbital Test Bed (2019–2021)

[4] Approximate performances are estimated from extrapolation of the performance claimed on the respective oscillator's datasheet.

2.2.2 Differences from the OCP-TAP Time Card

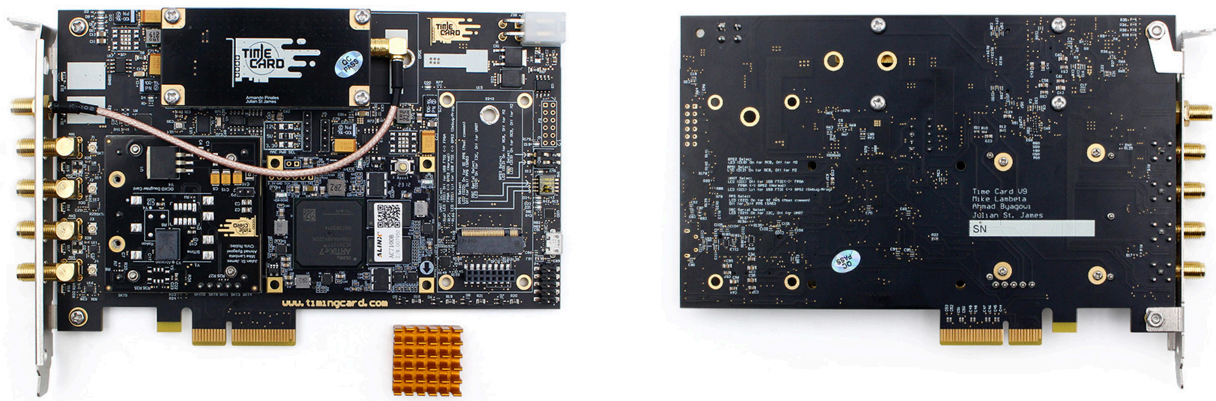

Figure 4. An integrated OCP-TAP Time Card v2 (Makerfabs).

While the OCP-TAP Time Card augments a host system that serves a network as an authoritative time appliance, the Space Time Card intentionally consolidates these roles to meet spacecraft autonomy requirements. Unlike many time appliances tailored for datacenter work, the Space Time Card must also withstand dynamic thermal and power conditions aboard the spacecraft. The Space Time Card omits features and components that are not appropriate or necessary for a spacecraft activities while remaining consistent with the OCP-TAP Time Card specification[5] and operates as an independent node and PTP Leader for the network. More experimentation is needed to qualify thermal and power sensitivities of the reference design.

Table 4. The baseline design only supports three modes: PTP Leader, Boundary Clock, and Passthru.

Input	Purpose	Space Time Card	
GNSS receivers	Up to 2 multi-band PPS + ToD	✅	Included as backup or to operate standalone from host
External PPS	Via SMA connector	🧪	Included for dev, omitted for flight.
External 10 MHz	Via SMA connector	🧪	Included for dev, omitted for flight.
IRIG-B	Time code input	❌	Not included, not used by the host.
DCF-77	Radio time signal	❌	Not included, not used by the host.
PTP	From network via PHC	✅	Included.
RTC	System clock as fallback	✅	Included.

[5] As of January 2026, IEEE Project Authorization Request P3335 is actively being drafted to standardize what constitutes a "Time Card," drawing on the OCP-TAP Time Card architecture (IEEE 2022).

Table 5. The baseline design integrates the Time Card with a minimal host for standalone operation, rather than being only a peripheral device.

Output	Purpose	Space Time Card
PPS	1 pulse per second, synchronized to adjustable clock	🧪 Included for dev, omitted for flight.
10 MHz	Reference frequency	🧪 Included for dev, omitted for flight.
PHC crossing	1-second crossing signal	❌ Not included. Host only interfaces over ethernet.
MAC PPS	Atomic clock's native PPS	❌ Not included. Host only interfaces over ethernet.
GNSS PPS	Pass-through from receiver	✅ Included as backup.
IRIG-B	Time code output	❌ Not included, not used by the host.
DCF-77	Radio time signal output	❌ Not included, not used by the host.
Configurable signal generator	Arbitrary frequency, phase, duration	❌ .Not included, not used by the host.

Table 6. The baseline design is optimized for operating as a standalone networked time appliance.

Interface	Purpose	Space Time Card
PCIe	Host communication, AXI bridge	❌ Not included, Host only interfaces with SoC. AXI is part of the internal architecture between FPGA and SoC.
4× SMA connectors	Configurable I/O	❌ Not included. Host only interfaces over ethernet.
Linux sysfs	Control/status interface that can be surfaced to Linux userspace via sysfs	❌ Not exposed to the host, only to internal SoC.
NMEA tty	Emulated serial port for legacy software	✅ Included for debugging.
PTP hardware timestamping	Network sync	✅ Included.

2.3 Separation of Concerns by Design

The Space Time Card lays foundations for infrastructure to grow incrementally and organically with the actual needs of a future Lunar ecosystem. Instead of trying to solve every problem today, the Space Time Card transforms the risk-cost-benefit profile of experimentation and iteration. This way the industry can pivot if there's a curveball without losing progress.

This design choice reflects a fundamental difference between terrestrial and spaceborne timing systems. In lunar and cislunar operations, continuous access to an external time authority cannot be assumed. The Space Time Card therefore operates as a self contained timing node that maintains time continuity locally, ingests external observables opportunistically, and exposes standardized interfaces for navigation and synchronization software.

When subsystems are tightly coupled, a change in one ripples through everything it touches, like drivers, commanding interfaces, new qualification cycles. Development slows to the pace of the slowest-moving dependency, and teams find themselves waiting on each other rather than making progress.

The Space Time Card manages to be both stable and hackable by separating functionality across well-defined boundaries. There are two primary interfaces that are central to the Space Time Card architecture:

1. A command and control interface to the host, including physical connections, drivers, and protocols that exchange data and power from the host to the device.

2. An interface between the device application and the timing package.

2.3.1 Distributing Precision Time

Precision time and PTP matured in highly networked terrestrial computing environments like telecom, automation, and high-frequency financial trading.[6] Time distribution has been the subject of decades of industry innovation, with mature technology routinely achieving sub-microsecond synchronization and, in controlled environments, sub-nanosecond precision between networked devices. The software for distributing time with this precision is likewise battle tested under extreme loads and edge-case conditions during nominal operation in terrestrial networks.

The Space Time Card package is Ethernet-addressible. All configuration, commanding, and time data exchanged across a local network with the host. The Time Card serves as the PTP Leader for the network and the host can follow standard practice for timing synchronization as a PTP Follower. Any device on the network can consume hardware-timestamped PTP packets as long as it is capable of being a PTP Follower. The host is free to use industry-standard software and patterns like White Rabbit to deliver precision time throughout the spacecraft. The Space Time Card thus inherits the robustness, performance, and feature set available to handle demanding

[6] NTP (Network Time Protocol) assumes symmetric latency between nodes, an assumption that breaks down below the microsecond level. PTP (Precision Time Protocol) uses hardware timestamping instead, enabling sub-microsecond synchronization on local networks and meeting the precision requirements described in this paper.

loads and conditions in modern datacenters.

1.2.3 Security Considerations and Trust Boundaries

Distributed timing systems face known attack vectors including time spoofing and coordinated drift injection. PTP networks face known attack vectors including grandmaster spoofing and delay manipulation, even with IEEE 1588-2019 security extensions (Schukat and Cortijo, 2021).

The current single-vehicle architecture places the security boundary at the host spacecraft: the Space Time Card trusts its internal oscillator ensemble, and the host is responsible for authenticating any external timing observables before passing them to the card.

For future inter-vehicle time exchange, Byzantine fault tolerance and cryptographic authentication of timing data are active research questions. The Epoch Project intends to operate within the LunaNet security framework for external communications.

2.3.2 Wide Compatibility and In-situ Upgrades

Flight software updates are routine spacecraft operations. Hardware changes after launch are not. The Space Time Card is designed around this asymmetry: network protocols and synchronization logic lives in software running on the card's SoC, and timing algorithms live in the card's FPGA. A mission that launches with basic PTP disciplining could adopt ensemble averaging or advanced holdover algorithms years later with no physical rework, just an upload.

This works because the interface to the host is deliberately minimal. The Space Time Card speaks PTP over Ethernet. The host is ignorant to whether there's a crystal oscillator or an atomic clock inside, vendor firmware or custom algorithms, one oscillator or an array of them. Integrating a Space Time Card into a new spacecraft requires attention to thermal management, power, and mechanical fit, but the data interface is already solved.

The same decoupling applies to hardware development for the Space Time Card itself. The SoC sees a stable oscillator through a daughter board rather than being soldered to the main PCB. Teams can swap oscillator technologies (OCXO or CSAC, single source or ensemble) with the same base card. The specifics related to a given stable oscillator are localized to the swappable daughter board. Multiple hardware variants with wholly dissimilar oscillators could operate with identical host software. The host would not know the difference.

A mother/daughter PCB architecture literally and figuratively decouples integration with the host from research and development of the timekeeping package. This architecture requires an additional hardware connection, more components, and more complexity. Objectively less ideal and unnecessary for a specialized system. The Space Time Card reference design accepts these drawbacks intentionally in exchange for being a flexible development platform.

The Space Time Card is not intended to be a monolithic instrument. It is a pattern for combining oscillators, discipline logic, relativistic transforms, and time transfer interfaces into a repeatable, interoperable subsystem. Missions may adopt the pattern with different performance and cost tradeoffs without redesigning their timing architecture from first principles.

Like Raspberry Pi and Arduino, the Space Time Card reference design is intentionally designed to be a “hackable platform” as well as a usable module in its own right. It lets users decide where (or where not) to spend their dollars, development time, and innovation. The interface remains stable so integrators, designers, and developers can build systems independently and have confidence that their systems will be compatible.

3. A New Epoch Starts Now

Why does this matter now? Lunar activity is accelerating across government, commercial, and academic actors. Programs led by organizations such as Intuitive Machines, Firefly Aerospace, ispace, Chinese lunar initiatives, ISRO's Chandrayaan missions, and European efforts are expanding the number of high value missions operating beyond Earth orbit. As this ecosystem grows, the need for precision timing, autonomy, and fault tolerance increases. At the same time, deploying dedicated lunar GNSS style infrastructure requires years of development and investments on the order of hundreds of millions of dollars, placing it out of reach for many near term missions. Meanwhile, the missions that would benefit most from advanced lunar timekeeping are launching now, before dedicated infrastructure exists to support them.

Onboard timekeeping provides an immediate and complementary capability. By carrying trusted time locally, spacecraft can reduce dependence on continuous Earth based synchronization and support precision navigation and coordination even when external services are unavailable or degraded. When multiple spacecraft each maintain validated onboard time, the system naturally becomes more resilient. Individual nodes can cross check, hold over, and recover without a single point of failure, enabling autonomy and graceful degradation.

Crucially, this capability does not require flagship class budgets. Mission appropriate precision can be achieved using relatively low cost components when they are paired with robust algorithms, disciplined holdover strategies, and rigorous validation. The performance ceiling is defined by architecture and implementation, not solely by the price of a single clock.

Open standards amplify this effect. A timing standard that is achievable by a university team is also achievable by major payload developers. Academic teams gain hands-on experience with deep space relevant hardware and algorithms, while industry benefits from a growing base of validated designs, data, and contributors. Whether a team contributes Space Time Card hardware, ensemble algorithms, or real time lunar frame transforms, each contribution advances the shared system. The underlying physics and techniques are well understood. The remaining challenge is implementation at scale, which is precisely the problem that motivated students and small teams are eager to solve.

3.1 Proposed Roadmap for Development

Epoch's development strategy is intentionally organized around three parallel pillars, each of which delivers standalone value and can be pursued independently. No single pillar is a prerequisite for progress in the others. This structure allows contributors from academia and industry to engage at different levels of hardware maturity, mission readiness, and research depth while still advancing a shared, interoperable system.

The roadmap emphasizes early validation, reuse of existing open infrastructure, and gradual expansion toward flight demonstrations rather than a single monolithic deployment. Answering research questions at this stage is more valuable than optimal performance.

Track 1: Space Time Card Hardware

The Space Time Card hardware effort focuses on adapting the OCP Time Card reference architecture for spaceflight use. This includes removal of datacenter specific interfaces, selection of space appropriate components, and packaging suitable for small spacecraft integration.

We have implemented a modified Allan deviation based weighting scheme on the Space Time Card V1, built in 2024-2025 by undergraduate engineering students at Rochester Institute of Technology, that combines short term stability from high performance crystal oscillators with long term stability from CSACs. Early bench testing shows that a two oscillator ensemble consisting of one Microchip SA.45s CSAC (predecessor of SA.65, current at the time of development) and one SiTime Stratum-3 Super TCXO maintains sub 50 nanosecond time deviation over 48 hour autonomous holdover periods under controlled laboratory conditions, compared to approximately 200 nanoseconds for the CSAC operating alone under similar conditions. This improvement is driven by regime specific weighting that suppresses short term phase noise while constraining long term drift, rather than by raw oscillator performance alone, and is not intended to imply equivalent results across arbitrary oscillator combinations or operating environments.

Initial work in 2026 centers on producing flight representative prototypes that preserve OCP compatibility while meeting spacecraft constraints. Hardware contributions may include oscillator daughter boards, power and thermal characterization, radiation tolerance studies, and interface validation.

Recommended milestones include bench level validation, followed by thermal vacuum and vibration testing beginning in 2027 for hardware that is ready for flight qualification. Hardware developed under this pillar is immediately useful as a platform for algorithm development and system integration, even prior to flight.

Track 2: On-Device Ensemble Algorithms

The ensemble algorithms development track focuses on improving timing stability and fault tolerance within a single node by combining multiple oscillators into a virtual clock. This work is primarily software and algorithmic and is intentionally decoupled from flight hardware availability. Internal ensembles of oscillators offer more stability than single-oscillator clocks at the expense of complexity. Whether an ensemble of low-cost oscillators can match the stability of a single higher-tier device remains an open research question.

Each Space Time Card operates as an independent timing node with validated onboard time. When multiple nodes are present, they may exchange timing observables through available communication links to estimate offsets, cross check behavior, and improve resilience. The architecture does not assume a single authoritative source and does not require continuous external synchronization.

Ensemble techniques can be developed, simulated, and validated using any OCP compatible Time Card or Space Time Card hardware, as well as through distributed laboratory deployments. Research questions include optimal weighting strategies, fault detection and isolation, and hybrid approaches that combine different oscillator classes for short term and long term stability.

Validated ensemble algorithms can be deployed remotely to networked Space Time Cards, enabling incremental upgrades without hardware redesign. Following sufficient validation, on-orbit deployment of ensemble algorithms is targeted no earlier than July 2027.

Track 3: Distributed Cislunar Networks

The distributed cislunar networks track is a forward-looking research track for developing and demonstrating concepts related to reconciling time across relativistic reference frames. The first objective is to implement relativistic time transformations, like those described by Ashby 2024, in real time on embedded hardware.. This pillar demonstrates that lunar rate corrections and frame transformations can be executed on low power flight class processors rather than being confined to ground systems.

Early demonstrations include Earth based experiments in which two precision clocks are maintained at different rates, synchronized to a common reference, to emulate Earth time and Lunar time side-by-side. These demonstrations validate numerical stability, resource usage, and integration with existing timekeeping infrastructure.

This development track depends on position knowledge but not on flight hardware. For Earth based demonstrations, fixed or simulated lunar ephemerides are sufficient. In flight or operational scenarios, ephemeris information may be provided by ground uplink or, eventually, inferred through the network itself.

3.2 Notional Development Timeline

The three development tracks proceed in parallel with different time horizons. Track 3 (lunar frame transforms) is the least constrained: Earth-based demonstrations using dual clocks at offset rates can begin immediately, independent of flight hardware or algorithm maturity. Track 2 (ensemble algorithms) can likewise advance on any OCP-compatible Time Card hardware available in the lab.

Track 1 (Space Time Card hardware) follows a more conventional qualification path. Bench-level validation of flight-representative prototypes is ongoing. Prototype hardware that is ready for qualification may enter thermal vacuum and vibration testing beginning in 2027. Following successful ground validation, on-orbit introduction of ensemble algorithms is targeted no earlier than July 2027.

Throughout all three tracks, supporting work includes shared test procedures, open documentation, and structured data-sharing practices. Early results from any track feed back into the others: oscillator characterization data from Track 1 informs weighting strategies in Track 2, and frame transform accuracy requirements from Track 3 set precision targets for both.

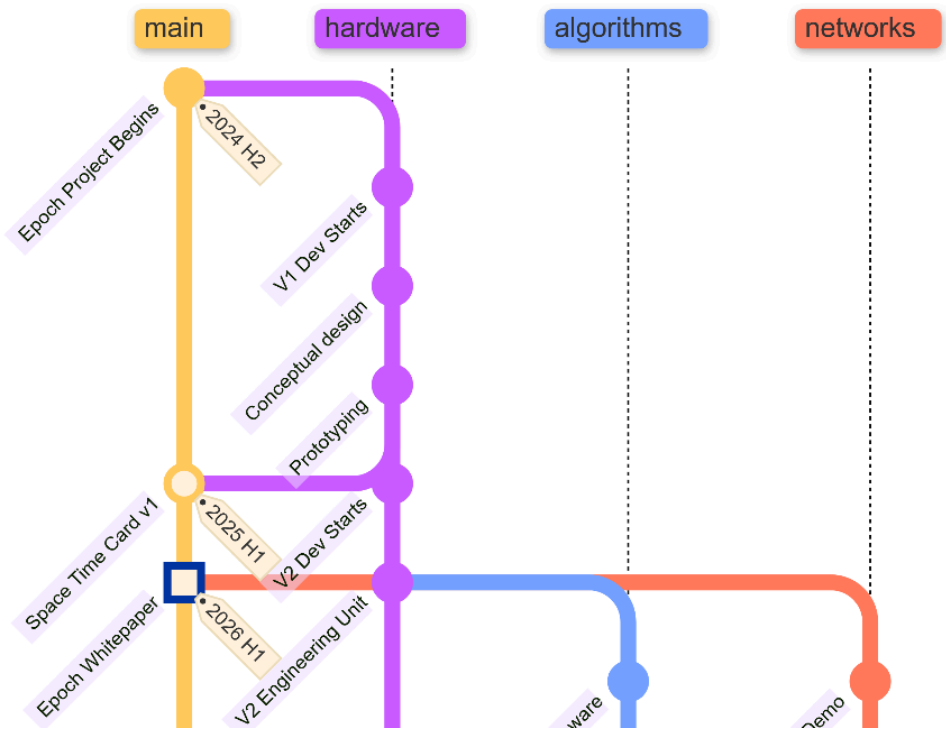
main
hardware
algorithms
networks
Epoch Project Begins
2024 H2
V1 Dev Starts
Conceptual design
Prototyping
Space Time Card v1
2025 H1
V2 Dev Starts
Epoch Whitepaper
2026 H1
V2 Engineering Unit

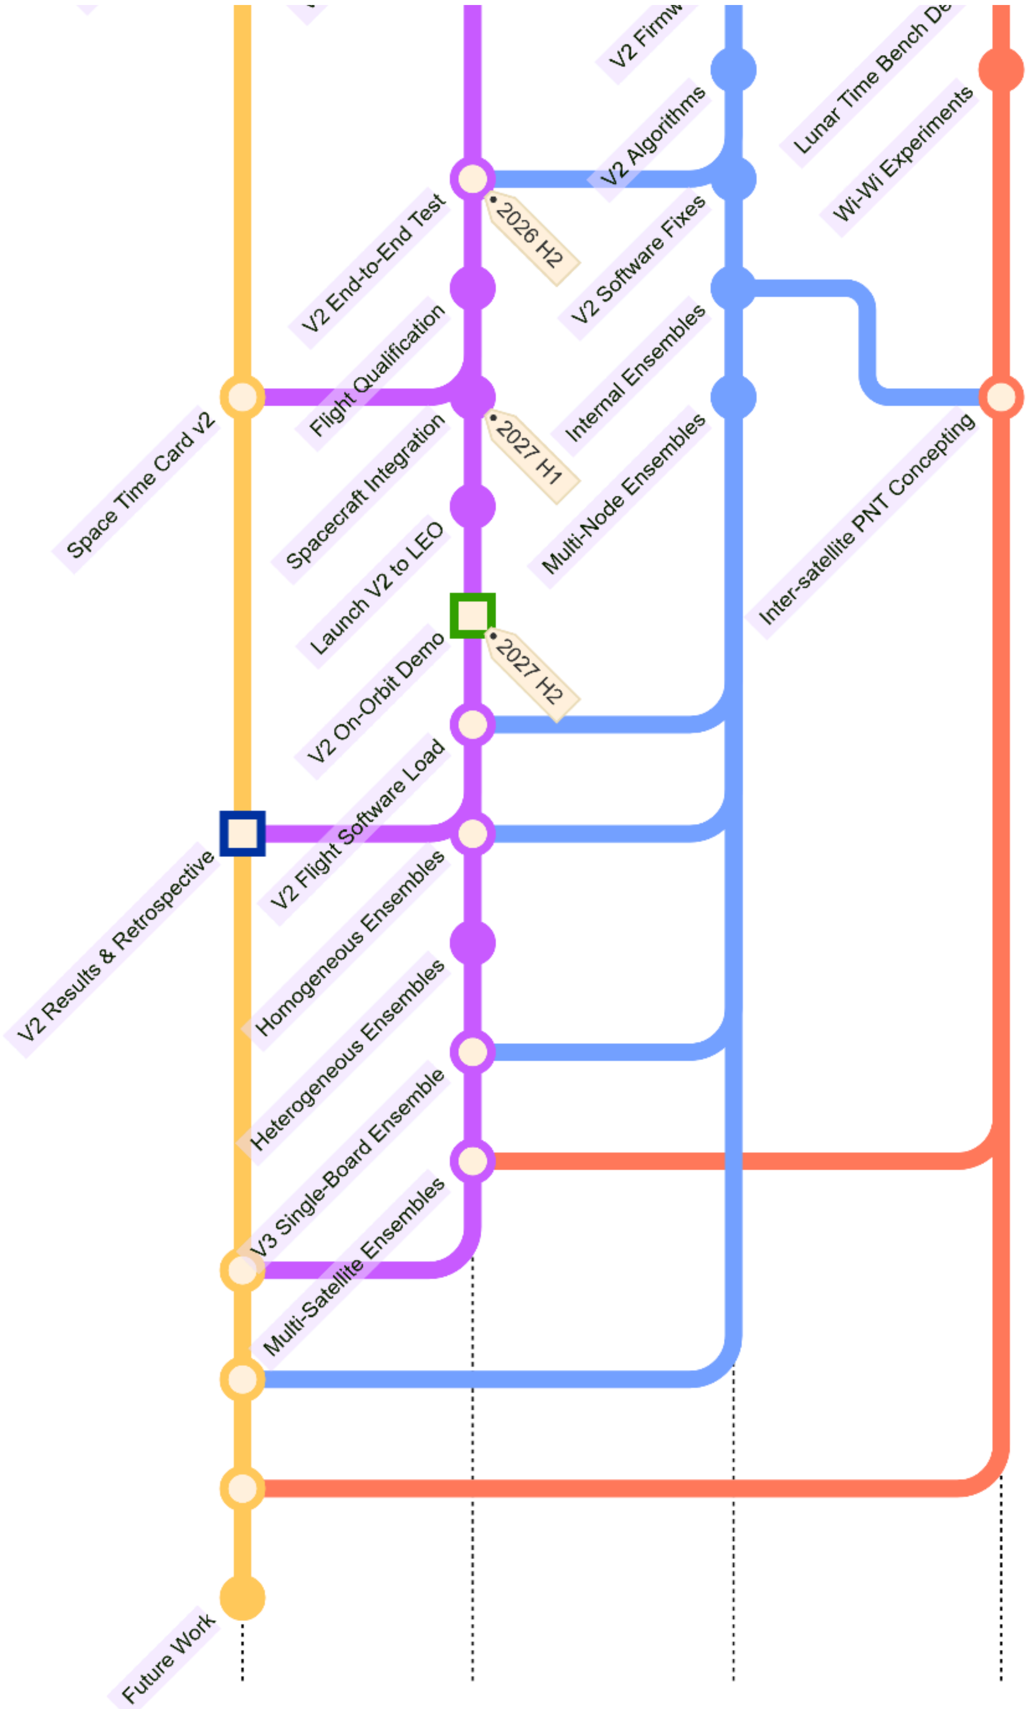

V2 Algorithms
Wi-Wi Experiments
V2 End-to-End Test
2026 H2
V2 Software Fixes
Flight Qualification
Internal Ensembles
Space Time Card v2
Spacecraft Integration
2027 H1
Multi-Node Ensembles
Inter-satellite PNT Concepting
Launch V2 to LEO
V2 On-Orbit Demo
2027 H2
V2 Flight Software Load
V2 Results & Retrospective
Homogeneous Ensembles
Heterogeneous Ensembles
V3 Single-Board Ensemble
Multi-Satellite Ensembles
Future Work

3.3 Let's Build the Lunar Timing Ecosystem Together

The Epoch Project invites contributions from university teams, researchers, industry engineers, and mission planners interested in advancing lunar timing infrastructure. Hardware design, ensemble algorithms, test procedures, validation data, and integration experiments all advance the shared ecosystem. Attribution and visibility will be provided to all contributors as the project evolves.

Contributions take many forms: hardware design, testing, and characterization; ensemble and fault-tolerant timing algorithms; embedded implementation of lunar time transforms; test procedures, validation data, and documentation; and integration experiments on ground or flight platforms. Each of these advances the shared ecosystem independently. A university team that characterizes oscillator performance in thermal vacuum produces data that benefits every future integrator. An algorithm researcher who improves ensemble weighting on a benchtop Time Card creates firmware that can be deployed to flight hardware without redesign.

In return, contributors gain access to an open reference design with real deep-space relevance, functional systems for deployment, and visibility within a growing community. The underlying physics and synchronization techniques are well understood. The remaining challenge is implementation at scale, which is precisely the kind of problem that motivated students and small teams are positioned to solve.

To get involved, visit the Epoch Project repository or contact the authors directly.

3.3.1 Open Research Questions

Several foundational questions remain open and represent concrete opportunities for collaboration. These are not prerequisites for the current development tracks—they are the questions that the tracks are designed to help answer.

Ensemble performance and composition. Can ensembles of many low-cost oscillators achieve stability comparable to a single higher-tier device? How does ensemble composition (homogeneous vs. heterogeneous, number of oscillators, oscillator class) affect short-term precision, long-term drift, and mission-lifetime resilience? These questions drive Track 2 development and can be investigated on benchtop hardware today.

Hardware in the space environment. Do miniature atomic clocks (MACs) function reliably in orbit? Is a space-qualified CSAC sufficient for LEO-based PNT applications? From GEO? Answering these requires flight data, which Track 1 is designed to produce.

Distributed time transfer. Is two-way carrier phase ranging (Wi-Wi) usable at orbital speeds and cislunar distances? Is there a viable peer-to-peer approach to a spaceborne PNT network, and if so, where are the boundaries and limiting sensitivities? These questions extend beyond the single-node scope of this paper and represent the next phase of research as the ecosystem matures.

OPEN LUNAR
FOUNDATION

References

Acknowledgements

The Epoch Project and this whitepaper would not have been possible without the generous contributions of time, expertise, and thoughtful critique from many individuals and organizations.

This paper benefited from detailed technical review by Vera Ko, Denis Reilly, Med Belhaj, Christopher Rabotin, and Dr. Dorian Patru Their feedback on network architecture, standards alignment, hardware selection, ensemble terminology, and security considerations substantially improved the accuracy and clarity of this work. Additional reviewers provided input on earlier drafts that helped clarify scope and positioning within the broader lunar infrastructure landscape.

The Space Time Card V1 was designed, built, and tested in 2024-2025 by student teams at Rochester Institute of Technology, whose work demonstrated the feasibility of our approach. The Open Lunar Foundation provided organizational support and the infrastructure investment philosophy that shaped the project's open-core approach. Olivia Linden contributed project support throughout. We are grateful for everyone's support

Any remaining errors or mischaracterizations are solely the responsibility of the authors.

Bibliography

Agasid, Elwood. 2024. "CAPSTONE: Cislunar Autonomous Positioning System Technology Operations and Navigation Experiment." NASA Technical Reports Server. https://ntrs.nasa.gov/citations/20240004138.

Ashby, Neil, Benjamin Bauch, and Jean-Luc Simon. 2024. "A Relativistic Framework to Establish Coordinate Time on the Moon and Beyond." NASA/TM-20240003144. NASA Technical Reports Server.

Burt, E. A., et al. 2021. "Demonstration of a Trapped Ion Atomic Clock in Space." *Nature* 595: 43–47.

Byagowi, Ahmad. 2021. "Open Sourcing a More Precise Time Appliance." Meta Engineering Blog, August 11, 2021.

Byagowi, Ahmad, et al. 2022. "Time Card and Open Time Server." In *Proceedings of the IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS).*

CERN. 2024. "CERN Launches the White Rabbit Collaboration." Press release, March 12, 2024.

CERN Open Science. 2026. "Open Hardware: White Rabbit." Accessed January 10, 2026. https://ohwr.org/project/white-rabbit.

ECSS. 2008. *ECSS-E-ST-10-04C: Space Environment.* European Cooperation for Space Standardization. https://ecss.nl/.

Flood, D., et al. 2023. "Ensemble Algorithms for Chip-Scale Atomic Clocks." *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 70 (8): 876–885.

Ginet, G. P., et al. 2013. "AE9, AP9 and SPM: New Models for Specifying the Trapped Energetic Particle and Space Plasma Environment." *Space Science Reviews* 179 (1–4): 579–615. https://doi.org/10.1007/s11214-013-9964-y.

Guo, J., et al. 2019. "Modeling of the Radiation Environment and Dose Assessment for a Lunar Base in a Lava Tube." *Space Weather* 17 (11): 1523–1532. https://doi.org/10.1029/2019SW002236.

Halekas, J. S., et al. 2008. "Lunar Prospector Observations of the Electrostatic Potential of the Lunar Surface and Its Response to Incident Currents." *Journal of Geophysical Research: Space Physics* 113 (A09102). https://doi.org/10.1029/2008JA013194.

IOAG and ICG (Interagency Operations Advisory Group and International Committee on GNSS). 2025. "Joint ICG-IOAG Multilateral Cislunar PNT Workshop." Vienna, February 11–13, 2025. https://ioag.org/meetings/cislunar/.

Manning, Catherine G. 2023. "The Interoperable Space Service Volume." NASA Space Communications and Navigation Program, September 25, 2023. https://www.nasa.gov/directorates/somd/space-communications-navigation-program/the-interoperable-ssv/.

Mazur, J. E., et al. 2011. "New Measurements of Total Ionizing Dose in the Lunar Environment." *Space Weather* 9 (S03002). https://doi.org/10.1029/2010SW000641.

Microchip Technology. 2025. *CSAC SA65 Chip-Scale Atomic Clock.* Datasheet DS00003478F.

Minow, J. I., et al. 2019. "The Lunar Surface Space Environment." In *Proceedings of the 2019 IEEE Aerospace Conference*, 1–8. https://doi.org/10.1109/AERO.2019.8741900.

NASA. 2017. *NASA-HDBK-4002A: Mitigating In-Space Charging Effects — A Guideline.* https://standards.nasa.gov/.

NASA. 2025. *LunaNet Interoperability Specification*, Version 5.0. NASA STI Program.

Normand, E. 1996. "Single Event Upset at Ground Level." *IEEE Transactions on Nuclear Science* 43 (6): 2742–2750. https://doi.org/10.1109/23.556861.

Open Compute Project. 2026. "Time Appliances Project." Accessed January 10, 2026. https://www.opencompute.org/projects/time-appliances-project.

Paige, D. A., et al. 2010. "Diviner Lunar Radiometer Observations of Cold Traps in the Moon's South Polar Region." *Science* 330 (6003): 479–482. https://doi.org/10.1126/science.1187726.

Petersen, E. L. 2011. "Single Event Analysis and Prediction." IEEE Nuclear and Space Radiation Effects Conference Short Course.

Schukat, M., and P. Cortijo. 2021. "Precision Time Protocol Attack Strategies and Their Resistance to Existing Security Extensions." *Cybersecurity* 4, article 12. https://doi.org/10.1186/s42400-021-00080-y.

Seubert, J., T. A. Ely, and J. Stuart. 2022. "Results of the Deep Space Atomic Clock Deep Space Navigation Analog Experiment." *Journal of Spacecraft and Rockets* 59 (6): 1914–1925. https://doi.org/10.2514/1.A35334.

Spence, H. E., et al. 2010. "CRaTER: The Cosmic Ray Telescope for the Effects of Radiation Experiment on the Lunar Reconnaissance Orbiter Mission." *Space Science Reviews* 150 (1–4): 243–284. https://doi.org/10.1007/s11214-009-9584-8.

U.S. Department of Transportation, Volpe National Transportation Systems Center. 2001. *Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System.* August 2001.

Williams, J.-P., et al. 2017. "The Global Surface Temperatures of the Moon as Measured by the Diviner Lunar Radiometer Experiment." *Icarus* 283: 300–325. https://doi.org/10.1016/j.icarus.2016.08.012.

Glossary of Terms and Abbreviations

Abbrev.	Term	Definition
	Aging, in oscillators	Slow, long term change in oscillator frequency over weeks to years, independent of short term noise.
ADEV	**Allan deviation**	A standard metric for frequency stability versus averaging time.
	Authority, time	A system whose observed time is recognized as the local reference.
	Autonomy	A spacecraft's ability to navigate, coordinate, and operate using onboard sensing and computation without continuous ground contact.
	AXI bridge	An internal bus interface used to connect time registers between an FPGA and SoC.
	Boundary clock	A PTP clock with multiple ports that synchronizes to one port and serves time on others, used to reduce timing fluctuations from intermediate network elements.
CAPS	**Cislunar Autonomous Positioning System**	An onboard navigation system demonstrated on the CAPSTONE mission, enabling peer-to-peer relative tracking between spacecraft without continuous ground-based support.
CAPSTONE	**Cislunar Autonomous Positioning System Technology Operations and Navigation Experiment**	A 12U CubeSat-class NASA mission launched June 28, 2022, that demonstrated autonomous navigation in a near rectilinear halo orbit around the Moon
CSAC	**Chip Scale Atomic Clock**	A compact atomic frequency reference offering high stability at low size, weight, and power, but with finite drift and aging that limit holdover without disciplining.
	Cislunar	The region of space encompassing the Earth Moon system, including

Abbrev.	Term	Definition
		lunar transfer, lunar orbit, and nearby operational volumes.
CERN	**Conseil Européen pour la Recherche Nucléaire**	The European Organization for Nuclear Research; developer of the White Rabbit timing protocol used as a precision synchronization reference in this paper.
CCSDS	**Consultative Committee for Space Data Systems**	An international standards body comprising major space agencies that develops and maintains data communication standards for spacecraft.
	Cross-discipline	A population of oscillators are coupled and use each other as their own reference.
	Crystal oscillator	An electronic oscillator that uses the mechanical resonance of a vibrating crystal to generate a precise frequency. The basis for most low-cost timing devices; stability degrades with temperature and aging without compensation.
DSAC	**Deep Space Atomic Clock**	A NASA technology demonstration showing exceptional onboard clock stability.
	Distributed timing	A model where multiple nodes maintain and validate time locally for resilience.
	Drift	Gradual divergence of a clock's time estimate from a reference.
	Ensemble	A virtual clock formed by combining multiple oscillators for improved stability and resilience.
	Ethernet	The network interface used for timing distribution and configuration. family of wired LAN physical and link layer standards commonly used to carry IP and timing protocols such as PTP.
ESA	**European Space Agency**	An intergovernmental organization of European member states responsible for coordinating Europe's space activities and programs.
FPGA	**Field Programmable Gate Array**	Reconfigurable logic used for precise timing and timestamping.
	Flight heritage	Operational evidence from flight that a design works in relevant environments, reducing integration risk for later missions.
GEO	**Geostationary Earth Orbit**	Earth orbit at approximately 35,786 km altitude where a satellite's orbital period matches Earth's rotation, appearing stationary from the ground.
GNSS	**Global Navigation Satellite System**	Satellite constellations that provide positioning and timing services in orbit, used here mainly as a contrast to the lunar environment.
GPS	**Global Positioning System**	A centralized satellite navigation and timing system used as both precedent and contrast.
	Ground based tracking	Navigation and timing operations performed primarily on Earth using tracking networks and post processing.
	Hackable platform	A deliberately modifiable reference system with stable external interfaces.
	In-situ upgrade	Post launch capability expansion via software rather than hardware changes.
IEEE	**Institute of Electrical and Electronics Engineers**	A professional association that develops and publishes technical standards widely adopted in electrical engineering, electronics, and related fields.

Abbrev.	Term	Definition
	Interoperability	The ability for different missions and systems to exchange timing data and observables using shared formats, interfaces, and conformance expectations.
	Light time delay	Signal propagation delay between spacecraft and ground that limits tight ground in the loop operations.
LEO	**Low Earth Orbit**	Earth orbit typically a few hundred to about two thousand kilometers altitude, often used as an initial flight environment for qualification and demonstrations.
LCT	**Lunar Coordinated Time**	Notional time scale for lunar operations issued by US Government in 2024. See also *Temps Coordonné Lunaire (TCL)*
	Lunar timekeeping	Maintaining a trusted time base for operations near the Moon, including modeling rate differences and supporting autonomous use onboard.
MEMS	**Microelectromechanical Systems**	A fabrication technology that integrates mechanical and electronic components at microscale; used in some oscillator designs as an alternative to quartz crystal resonators.
MAC	**Miniature Atomic Clock**	A small atomic frequency reference class that typically offers improved long term stability compared with CSAC class devices, at higher power, volume, or cost depending on implementation.
	Moonlight	An ESA initiative aiming to provide lunar communications and navigation services infrastructure.
	Mother and daughter board architecture	A hardware design separating base interfaces from oscillator specific integration. The base system is the *motherboard* and other circuit boards peripherally attached to it are *daughter* boards.
NASA	**National Aeronautics and Space Administration**	The United States federal agency responsible for civilian space exploration and aeronautics research.
NTP	**Network Time Protocol**	A protocol for synchronizing clocks over packet networks, typically achieving millisecond-level accuracy. Predates PTP and does not use hardware timestamping.
OCP	**Open Compute Project**	An industry consortium that develops and publishes open hardware and software designs for data center infrastructure, including the Time Card reference design.
OCP-TAP	**Open Compute Project Time Appliances Project**	A working group within OCP focused on open reference designs for precision timing hardware, including the Time Card and Open Time Server.
	Oscillator	A device producing a periodic signal used as a frequency reference.
OCXO	**Oven Controlled Crystal Oscillator**	A temperature stabilized crystal oscillator providing strong short term stability.
P2P	**Peer-to-peer**	Direct exchange of measurements between spacecraft without centralized authority.
PNT	**Positioning, Navigation, and Timing**	The combined capability set where trusted time supports navigation and coordination, and navigation context supports time transfer and frame transforms.
PTP	**Precision Time Protocol**	A network protocol for distributing time over packet networks, often

Abbrev.	Term	Definition
		using hardware timestamping to achieve high accuracy.
	PTP leader	The PTP clock selected to act as the time source for a domain, often called the grandmaster, distributing time to followers.
	Qualification, flight	Testing and evidence showing hardware meets mission environmental requirements.
RAFS	**Rubidium Atomic Frequency Standard**	A compact atomic clock using rubidium as its reference element, offering better long-term stability than CSACs at higher size, weight, and power.
	Range error	Positioning error expressed as an equivalent distance resulting from timing error.
	Rate offset	A persistent difference in clock rate between two time scales that accumulates into growing error if uncorrected.
	Resilience layer	Onboard capability that remains useful during outages or degraded conditions in external services or links.
RIT	**Rochester Institute of Technology**	A university in Rochester, New York; an RIT student team built the first Space Time Card prototype as part of a multidisciplinary senior design project sponsored by the Open Lunar Foundation.
SEU	**Single Event Upset**	A radiation induced bit or state upset in electronics.
	Space environment	Operational conditions including radiation, thermal extremes, and electromagnetic effects.
	Space Time Card	A reference implementation of the Epoch architecture providing onboard timing services.
	SpaceWire	A spacecraft communication standard developed by ESA for high-speed onboard data links, used as an alternative to Ethernet in some spacecraft architectures.
	Standard	A published definition of formats, behaviors, and tests enabling systems to work together.
SoC	**System on Chip**	An integrated processor managing configuration, telemetry, and networking.
τ	**Tau, averaging interval**	The averaging time interval over which frequency stability is measured in Allan deviation. A shorter τ captures short-term noise; a longer τ reveals drift and aging behavior.
Tech demo	**Technology demonstration**	A flight activity intended to validate new capabilities rather than support a primary mission.
TCXO	**Temperature-Compensated Crystal Oscillator**	A crystal oscillator with circuitry to correct for frequency variation across temperature, offering better stability than a bare crystal at modest cost and power.
TCL	**Temps Coordonné Lunaire**	A relativistic time scale used as a common reference for time tagging and comparison across lunar and cislunar nodes, accounting for relativistic rate differences. TCL is the relativistic coordinate time at the Moon's barycenter.
TT	**Terrestrial Time**	A standard Earth referenced time scale used as an input to transformations for lunar operations.

Abbrev.	Term	Definition
TVAC	**Thermal vacuum testing**	Ground testing combining vacuum and temperature extremes.
	Time appliance	A subsystem whose primary function is generating and distributing trusted time.
	Time base	The internal clock reference and time estimate used to timestamp events and interpret received signals.
TDEV	**Time deviation**	A timing stability metric expressed directly in units of time.
	Time registers	Software visible counters and timestamps representing timing state.
TID	**Total ionizing dose**	Accumulated radiation exposure degrading electronics over time.
TWTT	**Two Way Time Transfer**	A method using bidirectional exchanges to estimate clock offsets while reducing path delay uncertainty.
	Two-way tracking	A navigation method requiring signal turnaround, often tying observables to ground cadence.
	Vibration testing	Ground testing simulating launch and mechanical loads.
	Weighting scheme	A method assigning contributions from oscillators based on stability performance.
	Working group	A community forum for evolving interfaces, tests, and governance in a neutral manner.
	Workshop	A structured meeting convening representatives from multiple space agencies and international bodies to coordinate on a shared technical or policy problem.