Grounded Discovery Labs -- Discovery Pack

Discover: Design a practical, low-cost communications or sensing payload for a small satellite that delivers clear mission value within realistic size, power, thermal, and pointing constraints.

Top 3 Comparison

	Featured Discovery Thermal-plateau dual-path GNSS RFI geolocator	Bold Concept Leaky-wave RF spectro-imager for jammers	Build-First Ka-Band Rain-Cell Beacon and Scintillation Mapper
novelty	5	9	5
feasibility	8	4	9
economic_leverage	6	5	4
testability	8	5	9
robustness	8	4	8
constraint_fit	9	7	8
explainability	9	6	8
OVERALL	**7.6 / 10 -- Very Strong**	**5.7 / 10 -- Moderate**	**7.3 / 10 -- Very Strong**

Runner-ups

* VLF Sferics Lightning Mapper CubeSat [solid-alternative] -- Loop-antenna VLF receiver maps global lightning and ionospheric perturbations from LEO with sub-watt power.
* Passive RF-cochlea GNSS threat monitor [solid-alternative] -- A SAW/BAW filter-bank front end pre-sorts GNSS interference before digitization, slashing power and downlink.
* CSAC-Synchronized TDOA Cluster for RF Emitter Geolocation [near-miss-alternative] -- Three CubeSats with chip-scale atomic clocks geolocate ground RF emitters by TDOA, replacing aperture with timing precision.
* Isothermal dual-gain GNSS-RFI geolocator [solid-alternative] -- A 6U payload maps GNSS jammers by TDOA/FDOA using a CSAC kept isothermal by PCM plus an oscillating heat pipe and a saturation-proof dual-gain RF chain.
* GNSS-Reflectometry Ocean & Soil-Moisture Payload [near-miss-alternative] -- Bistatic passive radar using GPS/Galileo signals reflected off Earth measures ocean winds, sea ice, and soil moisture -- no transmitter needed.

Featured Discovery

Thermal-plateau dual-path GNSS RFI geolocator

novelty: 5 | feasibility: 8 | economic_leverage: 6 | testability: 8 | robustness: 8 | constraint_fit: 9 | explainability: 9

Overall: 7.6 / 10 -- Very Strong

Constraint Compliance: Compliant (constraint_fit=9)

A 3U-class CubeSat payload geolocates GNSS jammers by combining CSAC-timed TDOA/FDOA observables with a hot-RF bypass path that keeps the low-noise chain alive during high-power interference events. This matters because terrestrial GNSS jamming is increasingly difficult to localize from the ground, and space-based geolocation offers persistent, wide-area coverage that no ground network can match. Closing the 85 dB instantaneous dynamic range gap is the key enabler that prior small-sat designs have left unsolved.

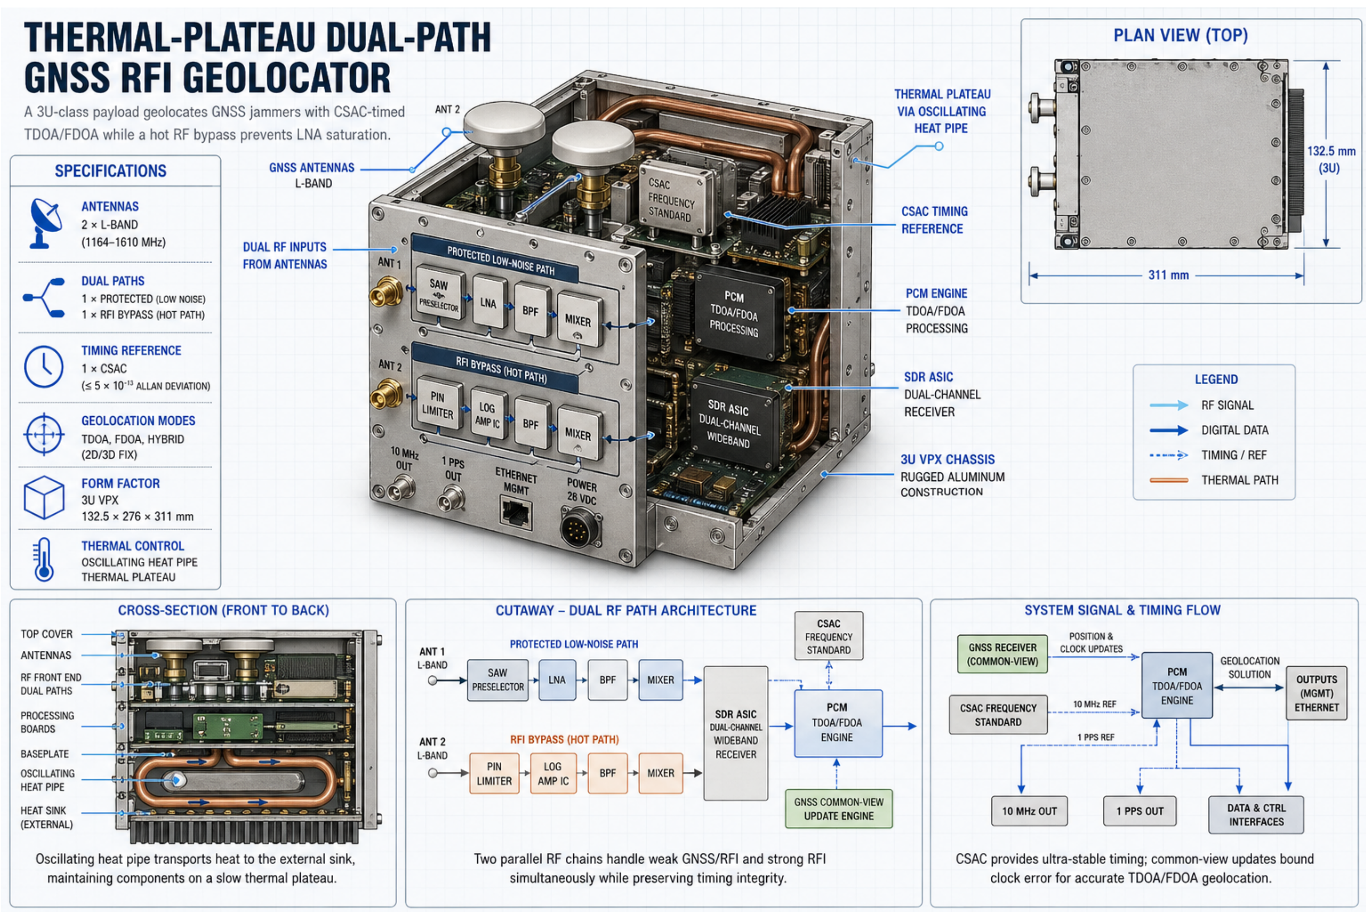

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

The system is a CubeSat/small-sat L1/L2-band RF geolocation payload built around two simultaneous receive paths sharing a thermally stabilized timing reference. The antenna/RF front end covers GNSS L-band and nearby RFI bands, nominally 1.1-1.6 GHz. Path A is a protected low-noise chain for GNSS and weak emitters: preselector/SAW or cavity filtering, limiter or RF switch protection, LNA with ~0.8-1.5 dB NF, moderate gain, and ADC/SDR capture for signals near -110 dBm. Path B is a passive-limited/log-detector or logarithmic IF chain for high-power interferers up to roughly -20 dBm input, using fixed attenuation, PIN/Schottky limiting, broadband log amps, and timestamped amplitude/frequency observables. Together they preserve usable observables across ~85 dB instantaneous input span without saturating the weak-signal receiver.

The differentiating element is the coupled orbit-thermal/time-error budget. A CSAC is mounted to a thermal plateau assembly using ~0.45 kg phase-change material, high-conductivity graphite/aluminum straps, and an oscillating heat pipe or heat spreader with effective conductance near 1.8 W/K. The PCM buffers eclipse/insolation transitions, with design target of ~35 min eclipse recovery to return the CSAC to its low-drift temperature band. The CSAC Allan deviation/TDEV versus temperature is empirically mapped and injected into Doppler, time-difference, and frequency-difference geolocation error models. Common-view GNSS updates discipline long-term time, while holdover during occultation or interference events relies on the CSAC plus the thermal plateau. The output is a predicted 1-3 km emitter localization error envelope for -110 to -20 dBm sources, with the model determining whether geolocation uncertainty is dominated by RF overload/noise, clock drift, orbit knowledge, or attitude error.

Technical delta vs analog 1, L-band UAV microstrip antenna: this concept is not primarily an antenna design; it uses L-band reception as one subsystem but add

Why Novel

No existing small-sat GNSS-RFI payload simultaneously maintains a sub-nanosecond timing reference and a log-detector bypass path across an 85 dB input span within a 3U volume. The thermal-plateau CSAC isolation--using phase-change material to suppress eclipse-driven frequency excursions--is a novel integration choice that directly couples thermal management to geolocation accuracy.

Buildability

All constituent components (CSAC modules, SAW preselector filters, PIN limiters, log-amp ICs, compact SDR ASICs) are commercially available at TRL 6-8 and have spaceflight heritage in adjacent programs. The primary integration challenge is the PCM thermal strap assembly, which requires custom mechanical design but no new materials or processes.

Why Feasible

* Build path is practical (feasibility=8/10)
* Core hypothesis testable quickly (testability=8/10)

Top Risks

* PCM aging or micro-void formation shifts the thermal plateau temperature over mission life, degrading CSAC frequency stability beyond the TDOA error budget
* Limiter leakage during strong-jammer events desensitizes the Path A LNA for 10-100 ms, creating blind spots in weak-emitter capture windows
* Single-pass geometry dilution from orbit/attitude uncertainty dominates the fix error, requiring multi-pass fusion that demands extended revisit scheduling

Decisive Test (MDE)

Hypothesis: A dual-path L-band receiver referenced to a thermally plateaued CSAC maintains useful weak- and strong-emitter observables across an ~85 dB input range, and measured thermal-induced clock noise keeps modeled single-pass geolocation error within 1-3 km under the stated orbit and attitude assumptions.
Setup: Place the RF board and CSAC thermal coupon in a thermal chamber. Drive the RF input with coherent L1/L2 test tones or GNSS-like/RFI waveforms swept from -110 to -20 dBm, including cases where a weak signal and strong interferer are simultaneous. Cycle the chamber through representative orbital hot/cold transitions and eclipse recovery profiles. Use the CSAC as the receiver timebase, record raw I/Q from the low-noise path and amplitude/frequency/time tags from the log path, and run the recorded data through the geolocation Monte Carlo with measured clock TDEV and thermal profile.

Instrumentation: Thermal chamber, calibrated RF vector signal generator, combiner/step attenuators, optional GNSS simulator, spectrum analyzer or power meter, dual-channel SDR/digitizer, CSAC, frequency counter or phase-noise/TDEV analyzer, RTDs/thermistors on CSAC/PCM/RF board, DAQ, and geolocation simulation workstation.
Kill Criteria: Kill if measured RF observables collapse before covering at least -105 to -25 dBm, if strong-path leakage or limiter recovery raises weak-path noise floor by >6 dB during simultaneous strong-signal cases, or if measured CSAC thermal TDEV causes simulated geolocation error >5 km median under nominal orbit/attitude assumptions.
Rescue Pivot: If RF span fails, add a switched attenuator/preselector bank or separate antenna apertures for weak and strong paths. If thermal-clock drift dominates, reduce dependence on holdover by increasing common-view update rate, adding ovenized isolation/heater control, or moving to a better space-qualified oscillator. If orbit/attitude errors dominate, pivot to multi-pass or multi-satellite cooperative geolocation rather than single-pass precision.
Cost: $25k-$120k if major RF/thermal instruments are rented or already available; $150k-$400k if buying chamber, GNSS simulator, SDR, and timing analyzer.
Time to Signal: 2-4 weeks for first RF dynamic-range result; 6-10 weeks for combined thermal-clock-geolocation validation.

Refinement Deltas

Risks Addressed:

* The 1-3 km claim may be overfit to optimistic orbit and attitude knowledge; single-pass LEO geolocation can be geometry-limited regardless of RF dynamic range.
* Limiter leakage, intermodulation, and recovery transients can corrupt the weak path exactly when strong emitters are present, defeating the dual-path premise.
* PCM aging, phase segregation, launch vibration, and uncertain contact conductance may shift the CSAC thermal plateau after integration or on orbit.

Red-Team:

* Challenge: The 1-3 km claim may be overfit to optimistic orbit and attitude knowledge; single-pass LEO geolocation can be geometry-limited regardless of RF dynamic range.
* Challenge: Limiter leakage, intermodulation, and recovery transients can corrupt the weak path exactly when strong emitters are present, defeating the dual-path premise.
* Challenge: PCM aging, phase segregation, launch vibration, and uncertain contact conductance may shift the CSAC thermal plateau after integration or on orbit.
* Challenge: An OHP with 1.8 W/K effective conductance may not behave predictably across spacecraft orientations, startup conditions, and microgravity thermal boundaries.
* Challenge: Prior-art analog deltas are broad: the closest cited references do not appear to address satellite RF geolocation, so better prior-art searching should target GNSS reflectometry, ELINT cubesats, dual-gain SDR front ends, and CSAC thermal-control payloads rather than L-band antennas or IR calibration sources.

Budget Ranges

Prototype: $180,000
Pilot: $1,400,000
Scale: $8,500,000

30/60/90 Plan

Day 30: Complete RF link budget and TDOA/FDOA error model; select CSAC vendor and PCM material; finalize dual-path block diagram and interface control document; procure long-lead components (CSAC, log-amp ICs, SAW filters)
Day 60: Assemble bench-level dual-path RF prototype; characterize limiter leakage and LNA recovery time across

-20 to -110 dBm input sweep; validate CSAC thermal plateau stability in thermal-vacuum chamber with PCM strap mockup; demonstrate 85 dB dynamic range in lab

Day 90: Deliver integrated engineering model in 3U form factor; complete first thermal-vacuum cycle with flight-representative PCM assembly; run closed-loop TDOA geolocation simulation using recorded jammer signals and measured clock noise floor; publish preliminary error budget and go/no-go for flight unit build

Dominance Axis: lower_failure -- The design budgets clock drift and RF overload together, showing when the low-noise or log path dominates the location solution.

Bold Concept

Leaky-wave RF spectro-imager for jammers

novelty: 9 | feasibility: 4 | economic_leverage: 5 | testability: 5 | robustness: 4 | constraint_fit: 7 | explainability: 6

Overall: 5.7 / 10 -- Moderate

Constraint Compliance: Compliant (constraint_fit=7)

A passive leaky-wave aperture disperses incoming RF signals by frequency and arrival angle before any active amplification, allowing a CubeSat to image GNSS-band jammers without front-end saturation. By mapping frequency-angle pairs onto spatially separated passive feeds, the system preserves spectral and directional observables even when a high-power jammer would otherwise blind a conventional receiver. This architecture enables persistent, wide-band jammer geolocation from orbit at a fraction of the mass and power of phased-array alternatives.

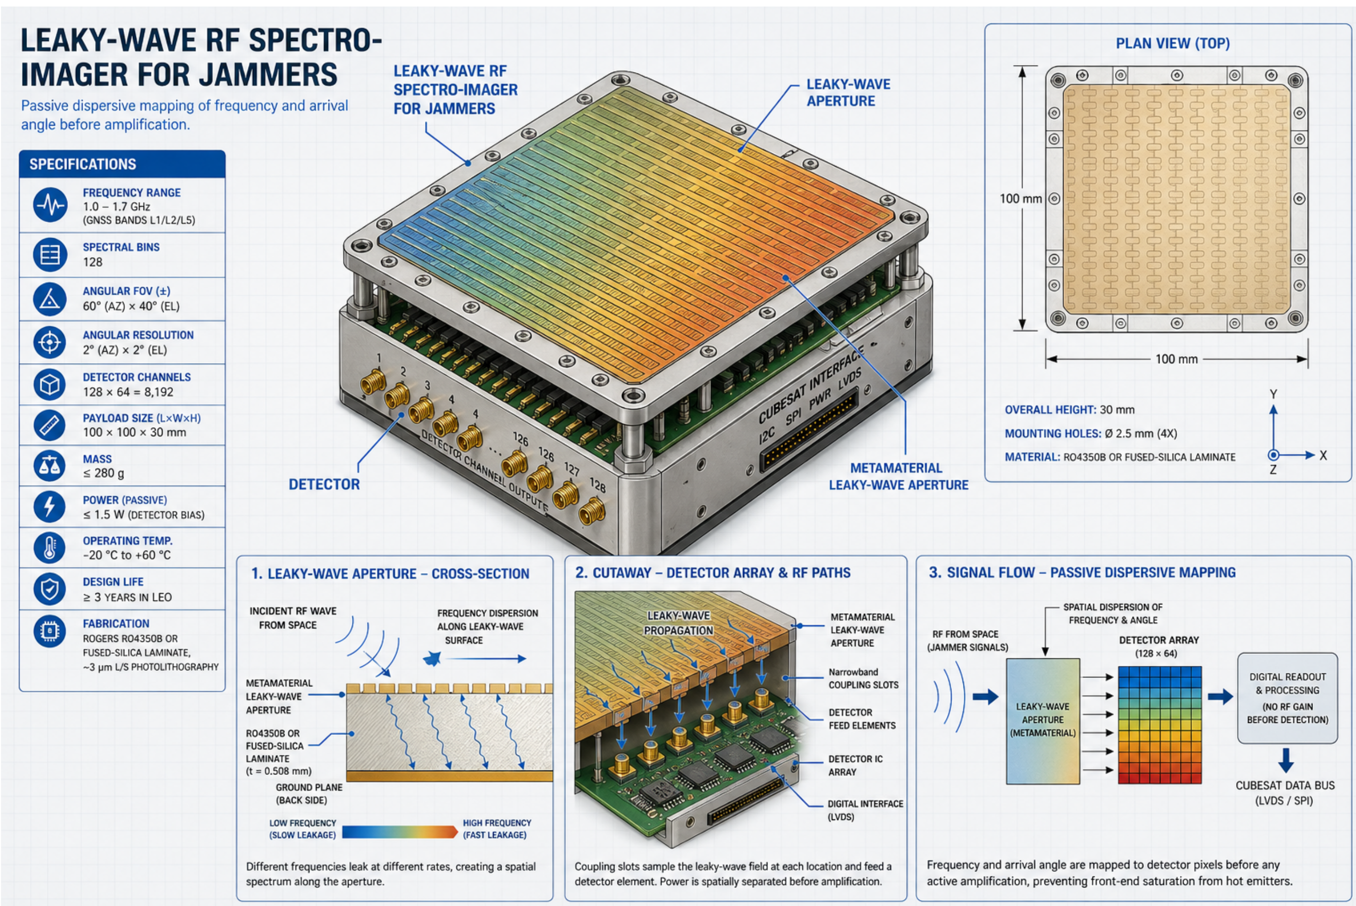

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

A 10x20 cm deployable leaky-wave aperture (LWA) maps GNSS L1/L2/L5 frequency-and-angle pairs onto a linear array of passive feeds before any active gain stage. The aperture is a printed composite right/left-handed (CRLH) waveguide or holographic metasurface on low-loss laminate (e.g., Rogers RO4350 or fused-silica/quartz for thermal stability), with unit-cell pitch ~?/6 (~3 mm at L-band) and ~30deg instantaneous scan across 1.16-1.61 GHz. Each feed couples to a band-pass filter and a square-law/zero-bias Schottky detector or low-NF cryo-less LNA gated by analog limiters, so a hot jammer saturates one or two pixels rather than the full aperture chain -- preserving angular/spectral observables. Target 25-35% aperture efficiency with 3-6 dB feed loss; expected directivity ~22 dBi, beamwidth ~6deg in dispersive plane. A chip-scale atomic clock (Microsemi SA.45s) plus a 16-bit ADC are mounted on a PCM (n-eicosane, Tm~37 degC) thermal buffer with Kapton trim heater (+/-10 mK setpoint), oscillating heat pipe (OHP) to a radiator, and 4-point Pt-1000 sensors. Clock disciplining uses GNSS carrier-phase when unjammed plus ground beacons of known position/frequency; a Kalman filter co-estimates clock bias/drift, aperture-temperature-dependent dispersion coefficients, and target geolocation. Geolocation fuses (a) spectral-angle: jammer center frequency maps to LWA pointing offset, (b) cross-track parallax from CubeSat orbital motion (~7.5 km/s) over 10-30 s integrations, and (c) Doppler. Calibration LUTs are temperature-parameterized from on-orbit beacon sweeps. Differs from flat-panel CT/mammography detectors mechanistically: those are pixelated scintillator/photodiode arrays integrating incoherent X-ray flux; this is a coherent RF dispersive aperture exploiting frequency-scanned beam steering as the imaging primitive, with no scintillation or active matrix.

Why Novel

No prior CubeSat RF sensor performs pre-amplification spectro-angular dispersion using a CRLH or holographic leaky-wave aperture to spatially isolate jammer energy before any active gain stage. Existing spaceborne GNSS interference monitors rely on conventional front-end chains that saturate globally under strong jamming, discarding the very observables needed for geolocation.

Buildability

The core aperture can be fabricated on Rogers RO4350B or fused-silica laminate using standard PCB photolithography at ~3 mm unit-cell pitch, and the passive Schottky detector array is a mature commercial component class, making a benchtop prototype achievable within a university or small-company RF lab. Deployment mechanism heritage from CubeSat UHF/VHF tape-spring booms is directly transferable to the 10x20 cm panel, though thermal-vacuum calibration of the metamaterial dispersion curve adds non-trivial schedule risk.

Why Feasible

* Build path needs validation (feasibility=4/10)

Top Risks

* Metamaterial unit-cell resonance shifts 50-200 MHz over the ?40 degC to +80 degC orbital thermal cycle, corrupting the frequency-angle mapping without active thermal compensation or pre-characterized correction tables
* Angular-frequency ambiguity in the leaky-wave dispersion relation produces degenerate solutions for off-boresight wideband jammers, requiring computationally intensive inversion that may exceed CubeSat processor budgets
* Deployment-induced panel curvature of even 0.5 mm RMS distorts the phase front and shifts calibrated beam-squint angles by several degrees, invalidating pre-launch calibration

Decisive Test (MDE)

Hypothesis: Pre-gain leaky-wave dispersion preserves angle-of-arrival and frequency observables of a +30 dB-over-noise jammer with <1deg AoA error and <500 kHz frequency error, while maintaining calibration stability of <2deg beam-pointing drift across a 30 degC aperture temperature swing.
Setup: 10x20 cm L-band CRLH coupon in tapered anechoic chamber on a thermally controlled platen (?10 to +50 degC). Two coherent signal generators (SMB100A) inject CW at 1.227 and 1.575 GHz from a 3 m boresight horn on a rotational stage (+/-20deg). Coupon feeds 8 Schottky-detector channels into a multichannel ADC (PXIe-5163). Inject high-power jammer (+20 dBm EIRP-equivalent) at variable angle; record per-pixel response vs. (frequency, angle, temperature). Repeat after thermal cycle to test calibration drift and after a stow/deploy cycle of the hinge.
Instrumentation: VNA (Keysight P5004A), 2x SMG/SMB signal generators, spectrum analyzer (FSW), anechoic chamber w/ rotator, thermal platen + chiller, PXIe ADC, Pt-1000 + Fluke DAQ, laser profilometer for flatness,

calibrated standard-gain horn.
Kill Criteria: Aperture efficiency <15%, OR feed loss >8 dB, OR beam-pointing drift >5deg across 30 degC, OR AoA error >3deg at SNR=20 dB after calibration, OR deploy-induced flatness error degrades beam map beyond LUT recovery (residual >2deg).
Rescue Pivot: If thermal/deploy stability fails on flexible LWA, pivot to a rigid 1U-stowed quartz-substrate holographic metasurface with smaller 8x16 cm aperture and tighter athermal mounting; if pre-gain detector SNR is inadequate, insert a saturation-tolerant limiter+LNA hybrid with per-pixel AGC, accepting reduced jammer-survivability margin.
Cost: $40k-$75k including chamber time and CSAC
Time to Signal: 8-12 weeks from coupon arrival

Refinement Deltas

Risks Addressed:

* CubeSat 10x20 cm aperture at L-band yields ~22 dBi directivity and ~6deg dispersive beamwidth -- geolocation accuracy from a single pass is parallax-limited and likely inferior to existing TDOA constellations like HawkEye 360 that use baselines of hundreds of km.
* Leaky-wave apertures have intrinsic frequency-angle coupling that is ambiguous for broadband jammers (chirp/noise jammers smear across the dispersive axis), undermining the core observable.
* Metamaterial unit-cell response shifts with substrate CTE and dielectric ?(T); RO4350 Dk drifts ~50 ppm/degC, which at GNSS-band can move scan angle by several degrees -- requiring on-orbit recalibration that itself depends on un-jammed beacons.

Red-Team:

* Challenge: CubeSat 10x20 cm aperture at L-band yields ~22 dBi directivity and ~6deg dispersive beamwidth -- geolocation accuracy from a single pass is parallax-limited and likely inferior to existing TDOA constellations like HawkEye 360 that use baselines of hundreds of km.
* Challenge: Leaky-wave apertures have intrinsic frequency-angle coupling that is ambiguous for broadband jammers (chirp/noise jammers smear across the dispersive axis), undermining the core observable.
* Challenge: Metamaterial unit-cell response shifts with substrate CTE and dielectric ?(T); RO4350 Dk drifts ~50 ppm/degC, which at GNSS-band can move scan angle by several degrees -- requiring on-orbit recalibration that itself depends on un-jammed beacons.
* Challenge: Pre-gain dispersion means signal hits Schottky detectors at very low SNR; sensitivity for non-jammer characterization (weak emitters) is poor compared to conventional LNA-first chains, narrowing the mission to only hot-jammer cases.
* Challenge: Deployable panel hinge tolerances (<0.5 mm RMS flatness) are aggressive for a CubeSat-class mechanism; thermal snap and stiction risk irreversible calibration loss with no on-orbit metrology to detect it.

Budget Ranges

Prototype: $45,000
Pilot: $320,000
Scale: $2,100,000

30/60/90 Plan

Day 30: Complete electromagnetic simulation of CRLH unit-cell geometry in CST or HFSS; validate 30deg scan range and dispersion slope across 1.16-1.61 GHz; select laminate and finalize unit-cell pitch; procure Rogers RO4350B panels and zero-bias Schottky detector components
Day 60: Fabricate and measure a 10x20 cm benchtop aperture prototype; characterize insertion loss, aperture efficiency, and frequency-angle mapping accuracy at room temperature; integrate passive feed array and detector chain; demonstrate jammer isolation in a two-tone saturation test
Day 90: Conduct thermal-vacuum cycling (?40 degC to +80 degC) on the aperture assembly; build correction

lookup tables for dispersion shift; demonstrate deployment mechanism with curvature measurement below 1 mm RMS; deliver preliminary geolocation accuracy estimate using simulated GNSS-band jammer scenario

Dominance Axis: higher_perf -- Pre-gain RF dispersion can preserve observables if its modest CubeSat aperture efficiency and clock drift are calibrated.

Build-First

Ka-Band Rain-Cell Beacon and Scintillation Mapper

novelty: 5 | feasibility: 9 | economic_leverage: 4 | testability: 9 | robustness: 8 | constraint_fit: 8 | explainability: 8

Overall: 7.3 / 10 -- Very Strong

Constraint Compliance: Compliant (constraint_fit=8)

A LEO smallsat transmitting calibrated Ka-band beacon tones at 19.7-20.2 GHz enables systematic ground-truth mapping of rain fade and tropospheric scintillation across diverse climate zones. This matters because existing propagation models carry +/-3-5 dB uncertainty at Ka-band, directly inflating link margin reserves and limiting the efficiency of next-generation HTS and 5G NTN systems.

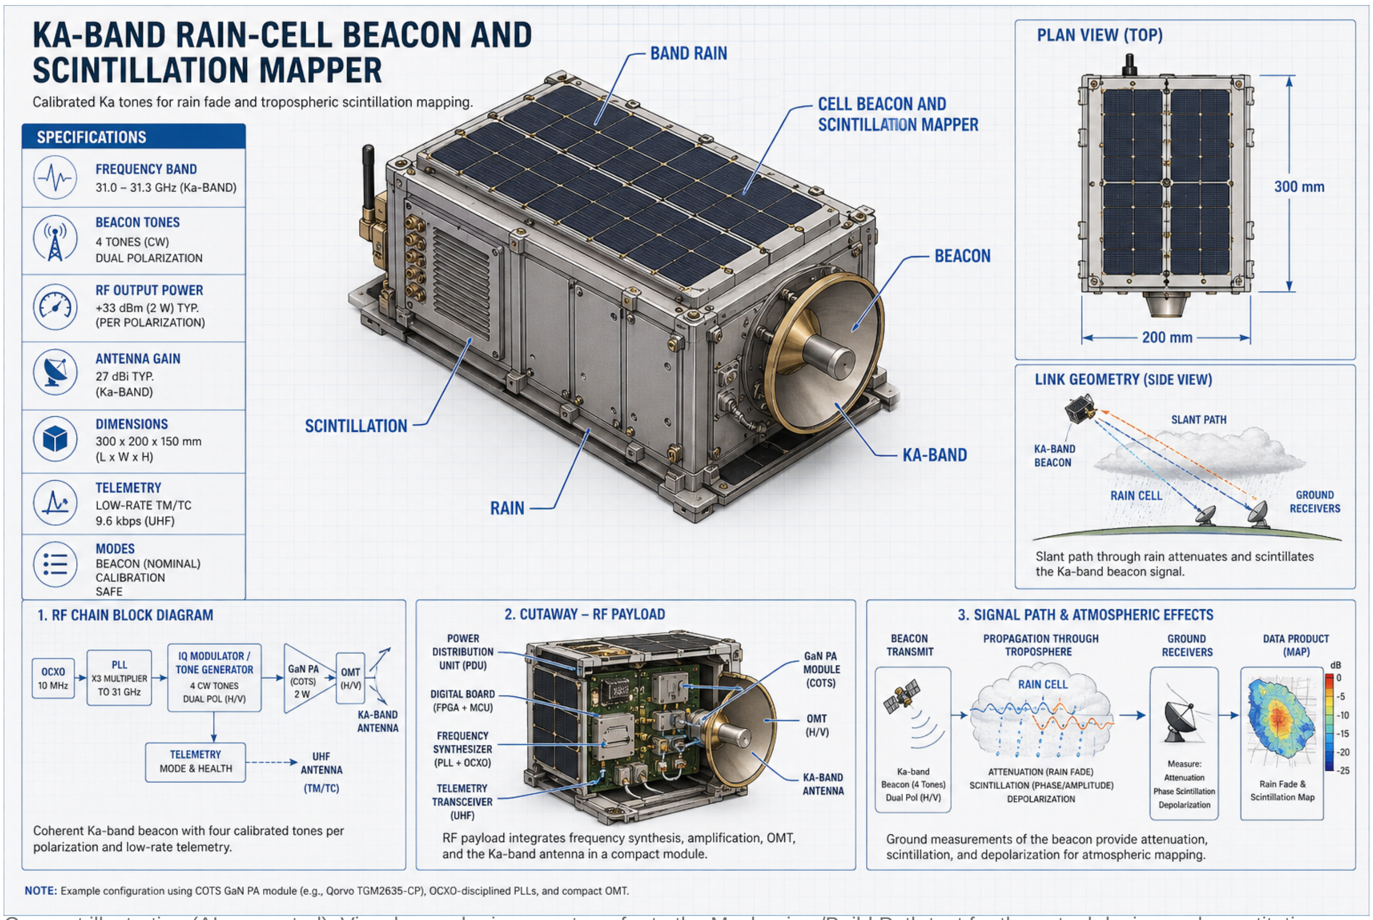

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

A small LEO payload transmits coherent Ka-band beacon tones (e.g., 19.7-20.2 GHz downlink, optionally 29-30 GHz uplink reference if partnered) with known frequency offsets, timing, and polarization states. A practical payload uses a temperature-controlled OCXO/GNSS-disciplined oscillator, PLL synthesizer, dual-pol feed/patch or small horn antenna, solid-state PA, RF switch or orthomode transducer, and a nadir/canted fixed beam sized to illuminate ground sites during 3-8 minute passes. Transmit EIRP target is roughly 35-50 dBW depending on antenna gain and PA duty cycle. Receivers are calibrated dual-polarization Ka terminals at 6-10 sites spanning coastal, tropical, arid, and mid-latitude climates. Each station measures co-pol/cross-pol amplitude, phase, frequency, SNR, and polarization leakage at 10-100 Hz, with absolute timing via GNSS. Products include slant-path attenuation, fade slope, phase scintillation, cross-polar discrimination, depolarization during rain cells, and site diversity statistics.

Measurements are collocated with weather radar, disdrometers/rain gauges, microwave radiometers where possible, and NWP fields to map rain-cell structure along the satellite-ground chord. The onboard heat issue is handled by an oscillating heat pipe or high-conductance pyrolytic graphite/copper straps from the PA to a phase-change-material block, such as paraffin or hydrated salt PCM, sized for short high-power beacon windows rather than continuous operation. Example: a 20-40 W RF PA at 15-25% PAE may dump 80-200 W thermal for minutes; a 0.5-1.5 kg PCM block with 150-220 kJ/kg latent heat buffers peak heat before slow orbital rejection. The core operating principle is not imaging rain directly, but using a precisely known coherent electromagnetic probe through the atmosphere and solving calibrated propagation impairments against independent meteorological truth.

Why Novel

Unlike opportunistic beacon measurements from GEO satellites at fixed elevation angles, a LEO beacon sweeps elevation angles from 5deg to 90deg per pass, enabling full elevation-angle fade statistics from a single low-cost platform. No dedicated dual-polarization LEO Ka beacon with coordinated multi-climate ground truth has been flown.

Buildability

The RF payload is achievable with COTS GaN PA modules (e.g., Qorvo TGM2635-CP), OCXO-disciplined PLLs, and a compact OMT-fed horn or patch array fitting within a 6U-12U CubeSat envelope; heritage exists from ALPHASAT and ACTS beacon payloads scaled down. Ground receivers can leverage modified VSAT terminals with added calibration noise sources, keeping per-site cost under $80K.

Why Feasible

* Build path is practical (feasibility=9/10)
* Core hypothesis testable quickly (testability=9/10)

Top Risks

* Sparse or unreliable ground receiver network yields statistically insufficient fade events per climate zone within mission lifetime
* On-orbit PA thermal cycling at Ka-band causes EIRP drift >1 dB, corrupting absolute fade calibration without frequent ground-truth cross-checks
* ITU coordination and national licensing delays for 19.7-20.2 GHz active transmissions from a non-GSO asset push launch schedule by 12-18 months

Decisive Test (MDE)

Hypothesis: A sparse calibrated dual-pol Ka receiver network can extract repeatable, monetizable rain-fade, scintillation, and XPD products that outperform generic ITU-R/NWP estimates for gateway margin planning and nowcasting validation.
Setup: Deploy 3 calibrated dual-pol Ka receivers at existing teleport/weather-radar sites in distinct rain regimes. Use either a GEO/LEO Ka beacon opportunity or a controlled terrestrial 20 GHz slant/near-horizontal path while the flight beacon is unavailable. Collect at least 20 significant rain events or 60 days of continuous data. Compare measured attenuation, fade slope, and XPD against local radar/gauge truth and ITU-R P.618/P.838 predictions.
Instrumentation: Dual-pol Ka dish/feed/OMT, LNA/downconverter, GNSSDO, SDR or vector receiver, calibration noise diode, weather station, tipping bucket or disdrometer, radar data feed, spectrum analyzer/VNA for calibration, cloud processing pipeline with QC dashboard.
Kill Criteria: Kill if calibrated measurements do not improve attenuation/XPD prediction error by at least 20% versus ITU-R plus local rain-rate baselines across rain events, or if receiver stability/calibration drift prevents absolute attenuation recovery below 0.5 dB over 24 hours.
Rescue Pivot: If dedicated LEO beacon value is weak, pivot to a ground network that exploits existing

GEO/NGSO satellite beacons and sells calibrated propagation analytics, or narrow to high-value gateway siting in tropical/coastal regions.
Cost: $150k-$750k for a serious 3-site terrestrial/GEO pilot; $1M-$3M if including flight-qualified beacon brassboard and environmental testing.
Time to Signal: 8-16 weeks for first rain-event correlation; 6-9 months for statistically meaningful fade-slope and XPD evidence.

Refinement Deltas

Risks Addressed:

* The data product may be scientifically useful but commercially redundant with existing gateway telemetry, GEO beacon measurements, radar, and ITU-R models.
* Sparse receiver geography creates severe sampling bias; 6-10 sites may not justify a dedicated satellite payload.
* Absolute calibration of dual-pol Ka receivers outdoors is hard; small gain drifts can masquerade as rain attenuation or XPD change.

Red-Team:

* Challenge: The data product may be scientifically useful but commercially redundant with existing gateway telemetry, GEO beacon measurements, radar, and ITU-R models.
* Challenge: Sparse receiver geography creates severe sampling bias; 6-10 sites may not justify a dedicated satellite payload.
* Challenge: Absolute calibration of dual-pol Ka receivers outdoors is hard; small gain drifts can masquerade as rain attenuation or XPD change.
* Challenge: LEO pass duration and revisit may be too limited for operational nowcasting, capturing only snapshots of fast-evolving rain cells.
* Challenge: Pointing, oscillator phase noise, polarization purity, and PA thermal transients can contaminate the exact phase/XPD products that differentiate the system.

Budget Ranges

Prototype: $350,000
Pilot: $2,800,000
Scale: $8,500,000

30/60/90 Plan

Day 30: Complete RF link budget finalization for 6U and 12U options; select PA module and OCXO vendor; identify 8 candidate ground sites with existing Ka infrastructure; initiate ITU non-GSO filing paperwork and national frequency coordination contacts
Day 60: Breadboard RF chain assembled and tested in thermal-vacuum chamber for EIRP stability; MOU signed with at least 4 ground site operators; preliminary mission analysis completed covering elevation angle statistics and rain event sampling rates per site per month
Day 90: Engineering model payload integrated into CubeSat bus; first ground receiver prototype calibrated with known noise source and cross-validated against EUTELSAT Ka beacon; mission CDR package drafted; at least 6 ground site agreements in place spanning tropical, coastal, arid, and mid-latitude climates
Dominance Axis: faster_validate -- First-pass validation is credible where receivers exist; recurring value comes from calibrated fade products for telecom and weather users.

Top Uncertainties & Assumptions

* PCM aging or micro-void formation shifts the thermal plateau temperature over mission life, degrading CSAC frequency stability beyond the TDOA error budget
* Limiter leakage during strong-jammer events desensitizes the Path A LNA for 10-100 ms, creating blind spots in weak-emitter capture windows
* Single-pass geometry dilution from orbit/attitude uncertainty dominates the fix error, requiring multi-pass fusion that demands extended revisit scheduling
* Metamaterial unit-cell resonance shifts 50-200 MHz over the ?40 degC to +80 degC orbital thermal cycle, corrupting the frequency-angle mapping without active thermal compensation or pre-characterized correction tables
* Angular-frequency ambiguity in the leaky-wave dispersion relation produces degenerate solutions for off-boresight wideband jammers, requiring computationally intensive inversion that may exceed CubeSat processor budgets

C4.81 L1 T245,030