OCS RESEARCH PAPER · PREPRINT · PAPER C (OBSERVATIONAL CAMPAIGN)

A Multi-Messenger Technosignature and Anomaly-Detection Campaign for Omega Centauri

Tim Swanson — The Omega Centauri Society / Post Oak Labs · tim@postoaklabs.com

Draft v1.0, June 2026 · companion to *The Macro Transcension Hypothesis* (Swanson 2026) · prepared for omegacentauri.me; intended for submission to arXiv and *Acta Astronautica* · HYPOTHESIS-AGNOSTIC CAMPAIGN — every program is conventional astrophysics first

⬇ PDF
⇔ Companion hypothesis paper (MTH)
⇔ Review of the inward-migration family
References
Decision tree

⚠ **EPISTEMIC STATUS.** This is the observational companion to the speculative Macro Transcension Hypothesis paper ("Paper A"). The campaign neither assumes nor argues for that hypothesis: every observation in it is independently justified by conventional astrophysics, and the technosignature analysis is a parallel pipeline on the same data. Sentences that depend on Paper A's specific predictions are explicitly labelled "Paper A dependence". Nothing here claims a detection of extraterrestrial intelligence.

ABSTRACT

Omega Centauri (NGC 5139), the most massive Galactic globular cluster and the probable stripped nucleus of an accreted dwarf galaxy, presents a unique conjunction of observational circumstances: the strongest current candidate for an intermediate-mass black hole (IMBH) in the Galaxy, anchored by seven stars moving above the local escape velocity (Häberle et al. 2024a); a formally unresolved factor-of-several tension between kinematic lower bounds ($\geq$8,200 $M_\odot$) and a pulsar-timing upper bound ($\lesssim$6,000 $M_\odot$; Bañares-Hernández et al. 2025); complete electromagnetic silence to the deepest radio and infrared limits ever placed on a globular cluster core (Mahida et al. 2026; Chen et al. 2025); a southern declination (−47.5°) optimal for the newest southern-hemisphere facilities; and

a 12-Gyr-old stellar population. No dedicated technosignature search of ω Cen has ever been conducted at any wavelength: the first dedicated globular-cluster SETI survey (Huang et al. 2026) could not observe it. We present a coordinated, hypothesis-agnostic, multi-messenger observational campaign that addresses both conventional astrophysical questions (IMBH reality, mass, and spin; cluster dynamics) and technosignature/anomaly hypotheses (including the inward-migration class surveyed in the companion paper) with the same data. The campaign comprises eight instrument-matched programs: (1) JWST infrared accretion and waste-heat limits; (2) deep radio imaging, narrowband SETI, and millisecond-pulsar timing with MeerKAT, transitioning to SKA-Mid; (3) astrometric monitoring with HST, Gaia DR4, the Nancy Grace Roman Space Telescope (launching August 2026), and ELT/MICADO; (4) a LISA extreme-mass-ratio-inspiral forecast; (5) a KM3NeT/ARCA neutrino burst pipeline exploiting ω Cen's up-going geometry from the Mediterranean; (6) Fermi-LAT archival analysis and CTAO-South follow-up; (7) optical/infrared time-domain and archival anomaly searches with Rubin/LSST and oMEGACat; and (8) supplementary archival channels (LIGO–Virgo–KAGRA continuous waves; Pierre Auger ultra-high-energy cosmic rays). For each program we state quantitative sensitivity estimates, time requests, decision thresholds, and explicit falsification criteria, including honest negative results: we show that direct astrometric acceleration detection of the fast stars is below 1σ at nominal parameters before ~2040, and route the astrometric science through photocentric-wander and reference-frame measurements instead. Total program cost is ≲ US$7M over 2026–2035, most of it archival analysis and piggyback observing on already-funded facilities; the decisive measurements (IMBH mass and spin via LISA) arrive as free by-products of planned mission science. The campaign is constructed so that every null result constrains conventional astrophysics, and every anomaly is adjudicated by at least two independent messengers.

Keywords: Omega Centauri · NGC 5139 · intermediate-mass black holes · technosignatures · SETI · multi-messenger astronomy · millisecond pulsars · gravitational waves · neutrino astronomy

Contents

1. Introduction

1.1 Why Omega Centauri, and why now

Three independent developments between 2024 and 2026 have transformed Omega Centauri (ω Centauri, NGC 5139) from one interesting globular cluster among many into arguably the single best-posed multi-messenger target in the Galaxy outside the Galactic Centre.

First, the IMBH candidacy became concrete. Häberle et al. (2024a), using a proper-motion catalogue of 1.4 million stars built from over 500 HST epochs (Häberle et al. 2024b), identified seven stars within 3 arcsec of the cluster centre moving faster than the local escape velocity — stars that require a bound central mass of ≥8,200 $M_\odot$ from their velocities alone, and ≥21,100 $M_\odot$ at 99 per cent confidence when acceleration limits are included. Independent N-body modelling finds a best-fitting in-situ-grown IMBH of ~4.7–5.1×10^4 $M_\odot$ (González Prieto et al. 2025). Against this, a joint analysis of stellar kinematics and millisecond-pulsar (MSP) timing places a 3σ *upper* limit of ~6×10^3 $M_\odot$ on any central point mass, favouring instead an extended ≈2–3×10^5 $M_\odot$ component of stellar remnants (Bañares-Hernández et al. 2025). The constraints are formally inconsistent under their stated assumptions. This is not an embarrassment but an opportunity: a factor-of-several disagreement between two mature measurement techniques, centred on the nearest IMBH candidate in the sky, is precisely the situation that focused observational campaigns exist to resolve.

Second, the electromagnetic silence became quantitative. The deepest radio observation ever made of a globular cluster — ~170 h with ATCA, reaching 1.1 μJy beam^{-1} rms — detects nothing at any proposed cluster centre, bounding the accretion efficiency of a putative IMBH at $\lesssim 4\times10^{-3}$ of the Bondi rate (Mahida et al. 2026), extending two decades of radio non-detections of globular-cluster IMBHs (Strader et al. 2012; Tremou et al. 2018). JWST NIRCam and MIRI photometry of the central field likewise finds no source with the spectral energy distribution of an accreting IMBH (Chen et al. 2025). Whatever sits at the centre of ω Cen, it is dark to a depth that is itself a scientific result demanding explanation.

Third, the instrument landscape shifted decisively in the campaign's favour. Within an eighteen-month window centred on this writing: the Vera C. Rubin Observatory began the Legacy Survey of Space and Time (LSST) in early 2026 (Ivezić et al. 2019); the Nancy Grace Roman Space Telescope completed construction and is scheduled for launch on 2026 August 30 (Sanderson et al. 2019); Gaia Data Release 4, with a 5.5-year astrometric baseline, is scheduled for 2026 December 2 (Gaia Collaboration 2016); KM3NeT/ARCA reached 51 of 230 planned detection units with real-time alerts entering commissioning (Adrián-Martínez et al. 2016; KM3NeT Collaboration 2025); the first telescopes of CTAO-South are arriving at Paranal (CTA Consortium 2019); SKA-Mid is assembling toward 2029 science verification (Braun et al. 2019); ELT first light is scheduled for 2029 (Davies et al. 2021); and LISA, formally adopted by ESA in January 2024, is in hardware development for a 2035 launch (Colpi et al. 2024). Nearly every facility on this list is southern-hemisphere or all-sky; ω Cen, at declination −47.5°, sits in the sweet spot of all of them. A campaign organized now can ride this entire wave at marginal cost.

1.2 The technosignature rationale, stated honestly

This paper is the observational companion to a speculative hypothesis paper (Swanson 2026, hereafter Paper A), which argues that rapidly spinning massive black holes in dense old stellar systems are thermodynamic attractors for hypothetical computation-optimizing civilizations, and that ω Cen is the most accessible system satisfying the resulting selection criteria. We neither assume nor argue for that hypothesis here. The campaign is designed to be *hypothesis-agnostic* in the specific sense that every observation in it is independently justified by conventional astrophysics: IMBH demographics (Greene et al. 2020), cluster dynamics, pulsar timing, accretion physics at the lowest Eddington ratios, and EMRI astrophysics (Amaro-Seoane 2018). The technosignature dimension

adds analysis channels — narrowband drift searches, burst-coincidence triggers, anomaly statistics — to data that would be taken anyway, in the cost-effective tradition of commensal SETI (Tarter 2001; Wright et al. 2022; Lacki & DiKerby 2025).

Three considerations nevertheless justify making the technosignature dimension explicit rather than incidental. First, globular clusters are almost entirely unexplored SETI territory: the first dedicated globular-cluster survey, a FAST pilot of five northern clusters, was published only this year and could not observe ω Cen from FAST's latitude (Huang et al. 2026). Second, ω Cen is a privileged target on entirely generic arguments — 10^7 stars of age 12.08 ± 0.01 Gyr (Clontz et al. 2024) in a beam a few arcminutes across, a possible IMBH, and (as the probable nucleus of an accreted dwarf; Hilker & Richtler 2000; Ibata et al. 2019) an independent galactic chemical-evolution history — so that any ETI prior that weights stellar age and number density at all concentrates sharply here. Third, the high-energy technosignature channels (neutrino bursts from hypothetical black-hole-based computation; Dvali & Osmanov 2023) happen to be testable at ω Cen essentially for free, because the cluster's declination makes it an up-going source for the Mediterranean neutrino telescopes. Where Paper A's specific predictions sharpen a test, we say so explicitly and label the dependence; readers uninterested in that hypothesis may strike those sentences without loss to the campaign.

1.3 Design principles

Five principles govern the campaign design, and we state them up front because they discipline everything that follows.

(i) Dual-use data. Every observing program must produce publishable conventional astrophysics on its own; the technosignature analysis is a parallel pipeline on the same photons (or neutrinos, or strain data), never a dedicated expenditure that a null result would waste.

(ii) Honest sensitivity accounting. Where a measurement cannot work at nominal parameters, we say so quantitatively rather than gesturing at it hopefully. Section 5 demonstrates this standard: the direct acceleration test on the fast stars — superficially the most natural follow-up to Häberle et al. (2024a) — is below 1σ before ~2040 for plausible masses, and we restructure the astrometric program accordingly.

(iii) Two-messenger adjudication. No anomaly claim survives on one channel. Every positive trigger (a neutrino multiplet, a narrowband line, an infrared transient) must specify in advance which

independent channel confirms or kills it, with combined significance assessed by pre-registered methods (Section 11).

(iv) Pre-registered thresholds. Detection and falsification thresholds are stated numerically before the data arrive (Tables 3 and 4), in the tradition that a test postponed until after the data are seen is not a test.

(v) Parasitism on funded science. The decisive measurements — IMBH reality, mass, and spin — will be made by LISA as part of its core science program whether or not anyone organizes a campaign (Colpi et al. 2024; Babak et al. 2017). The campaign's job is to ensure that when those numbers arrive, the contextual data (timing, astrometry, electromagnetic limits, high-energy monitoring) already exist to interpret them.

1.4 Structure of this paper

Section 2 reviews the target and the present constraint landscape. Sections 3–10 present the eight programs, each with science case, technical implementation, sensitivity estimates, and decision criteria. Section 11 describes cross-program coordination, alert protocols, and joint statistics. Section 12 presents costing and timeline. Section 13 assembles the decision tree through 2040. Section 14 concludes.

2. The target and the constraint landscape

2.1 Adopted parameters

Table 1 collects the parameters adopted throughout. We use the oMEGACat kinematic distance of 5.49 ± 0.06 kpc (Häberle et al. 2025), consistent within systematics with the Gaia EDR3 parallax distance (Soltis et al. 2021) and the combined-catalogue value (Baumgardt & Vasiliev 2021). At this distance 1" subtends 0.0266 pc, so the entire fast-star region ($r < 3"$) fits within 0.08 pc, and the cluster core ($r_c \approx 2.4'$) within ~4 pc.

Table 1. Adopted parameters for ω Centauri (NGC 5139).

QUANTITY	VALUE	SOURCE

QUANTITY	VALUE	SOURCE
Distance	5.49 ± 0.06 kpc	Häberle et al. (2025)
Cluster mass	≈ 4×10^{6} $M_{\odot}$	Baumgardt & Hilker (2018)
Stellar count	~10^{7}	Baumgardt & Hilker (2018)
Mean stellar age	12.08 ± 0.01 Gyr (spread 0.75 Gyr)	Clontz et al. (2024)
Origin	stripped dwarf-galaxy nucleus	Hilker & Richtler (2000); Ibata et al. (2019)
Position (J2000)	RA $13^{h}26^{m}47^{s}$, Dec −47°28'46"	Harris (1996)
Known millisecond pulsars	18	Dai et al. (2020); Chen et al. (2023)
Central γ-ray source	4FGL J1326.7−4729 (MSP ensemble)	Dai et al. (2020, 2023)
Angular scale	1" = 0.0266 pc	—

2.2 The mass tension

Table 2 and Figure 1 summarize the central-mass constraints. Two features matter for campaign design. First, the disagreement is not marginal: the velocity-only kinematic lower bound (8,200 $M_{\odot}$) exceeds the pulsar-timing point-mass upper bound (6,000 $M_{\odot}$, 3σ) outright, and the acceleration-informed lower bound (21,100 $M_{\odot}$) exceeds it by a factor of 3.5. At least one analysis is wrong, or incomplete, in an instructive way — for instance through the radial distribution assumed for the MSPs, the foreground/membership treatment of the fast stars, or the possibility that the central mass is genuinely extended (Bañares-Hernández et al. 2025; Zocchi et al. 2019; Breen & Heggie 2013). Second, the tension is *resolvable on a known schedule*: more pulsars and longer timing baselines tighten the upper bound as roughly $1/\sqrt{(N_{psr}\,T^{2})}$ (Section 4); Roman and Gaia DR4 sharpen the astrometric frame within two years (Section 5); and LISA, if a compact-object inspiral is caught, measures the mass to ~0.1 per cent (Babak et al. 2017).

Table 2. Current constraints on the central dark mass of ω Cen (cf. Paper A, Table 2).

CONSTRAINT	VALUE	METHOD / SOURCE

CONSTRAINT	VALUE	METHOD / SOURCE
Lower bound (velocities only)	$\geq$ 8,200 $M_\odot$	7 fast stars, HST proper motions (Häberle et al. 2024a)
Lower bound (with accelerations)	$\geq$ 21,100 $M_\odot$ (99%)	same stars, acceleration limits (Häberle et al. 2024a)
N-body growth models	~4.7–5.1×10^4 $M_\odot$	cluster-evolution simulations (González Prieto et al. 2025)
Upper bound (point mass)	< 6×10^3 $M_\odot$ (3σ)	stellar kinematics + MSP timing (Bañares-Hernández et al. 2025)
Favoured alternative	extended 2–3×10^5 $M_\odot$	same analysis (Bañares-Hernández et al. 2025)
Historical upper limits	$\lesssim 10^4$ $M_\odot$ (model-dep.)	HST proper motions (van der Marel & Anderson 2010); see also Noyola et al. (2008); Zocchi et al. (2019)

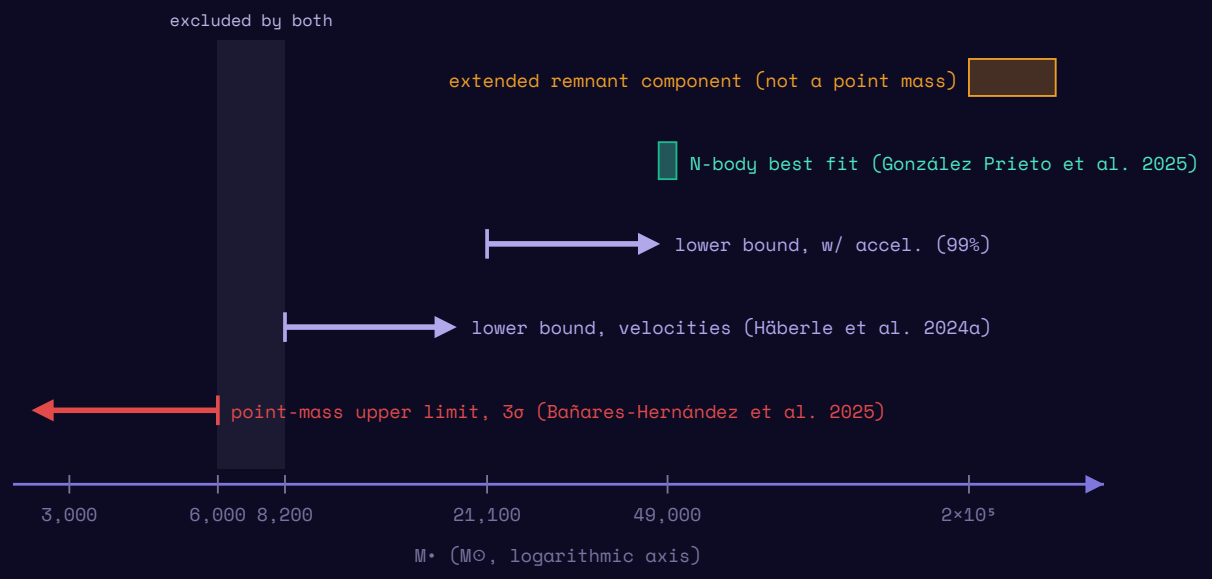

Figure 1. The central-mass constraint landscape (logarithmic mass axis). The pulsar-timing point-mass upper limit and the fast-star kinematic lower bounds leave *no* allowed point mass: at least one constraint must yield. The campaign is structured to determine which, on a known schedule, using instruments that share no dominant systematics.

2.3 Three hypotheses, fixed in advance

Following Paper A we carry three competing hypotheses through every program:

> **H_0 (gas-starved IMBH or near-IMBH).** A genuine central point mass in the ~10^4 $M_\odot$ range exists and is silent for the ordinary reason: a relaxed, gas-poor, 12-Gyr-old cluster supplies essentially nothing to accrete. Default and most parsimonious if the kinematic bounds hold.

> **H_1 (extended remnant component).** No IMBH; the central mass is a spatially extended cluster of stellar remnants (Bañares-Hernández et al. 2025; Breen & Heggie 2013; Zocchi et al. 2019). Default and most parsimonious if the timing bound holds.

> **H_2 (anomalous/managed system).** A central IMBH whose observational properties are shaped by technology, per Paper A. Carries a low prior; becomes interesting only through conjunctions that H_0 and H_1 jointly fail to explain (e.g. near-extremal spin on a gas-starved hole; repeated coincident multi-messenger bursts; statistically anomalous depletion of the loosely bound core population).

The campaign's purpose is *not* to hunt for H_2; it is to measure the system well enough that H_0 and H_1 are decided on their merits — with H_2's distinctive residues either appearing in the data or being retired by them.

3. Program 1: JWST infrared accretion and waste-heat limits

3.1 Science case

For any central mass $M_\bullet$, the Eddington luminosity is $L_{Edd} \simeq 1.26\times10^{31}\ (M_\bullet/M_\odot)$ W: 1.0×10^{35} W at 8,200 $M_\odot$, 2.5×10^{35} W at 2×10^4 $M_\odot$, 6.2×10^{35} W at 4.9×10^4 $M_\odot$. The existing non-detections (Mahida et al. 2026; Chen et al. 2025) already establish that any IMBH radiates at $\lesssim10^{-9}$ of Eddington — among the most extreme quiescence levels measured for any black hole — but the published JWST analysis is an archival by-product, not an optimized deep program. A purpose-designed campaign improves the limits by an order of magnitude and adds the two dimensions the archival data lack: *variability* (accretion at these rates is expected to flicker; a managed or artificial

suppression of accretion, the H_2 residue, is not) and *waste heat* (a mid-infrared excess over the stellar population model is the classic Dyson-type technosignature; Dyson 1960; Wright et al. 2014; Hsiao et al. 2021 — and at ω Cen the relevant solid angle is a few arcseconds, not a galaxy).

3.2 Implementation

Deep imaging (new observations). NIRCam F200W/F356W/F444W at 20 ks per filter plus MIRI F770W/F1000W at 15 ks per filter, in two epochs separated by ~6 months: ~65 h total. In the crowded core the limits are confusion- rather than photon-limited; with empirical-PSF subtraction anchored to the oMEGACat astrometric catalogue (Häberle et al. 2024b; Nitschai et al. 2023), simulated point-source recovery reaches ~5–10 nJy (3σ) at F444W and ~30–70 nJy at F770W/F1000W after a ~3× crowding penalty. At 5.49 kpc, 10 nJy at F444W corresponds to a monochromatic $\nu L_\nu \approx 2.5\times10^{28}$ erg s^{-1}; for plausible quiescent-accretion SEDs the bolometric limit is of order 10^{29}–10^{30} erg s^{-1}, i.e. Eddington ratios of 10^{-13}–10^{-12} at 2×10^4 $M_\odot$ — probing a regime in which even Sgr A* would be conspicuous (Event Horizon Telescope Collaboration 2022).

Two-epoch variability. Differential photometry between epochs is robust to the static crowding systematics that dominate the absolute limits; variability at the ≥20–30 per cent level for any source above ~30 nJy flags an accretion candidate and triggers the coordination protocol of Section 11.

Waste-heat analysis (archival + new). The 10–30 μm field photometry is fitted against the synthetic stellar-population SED constructed from the oMEGACat spectroscopic catalogue (Nitschai et al. 2023); the technosignature statistic is a spatially coherent excess in the central 0.1 pc exceeding 3σ over the population model. The same fit yields conventional science: the most precise mid-IR census of the post-main-sequence population in any globular cluster.

Spectroscopic environment test (Cycle 6+, contingent). If astrometry or timing strengthens the IMBH case, a NIRSpec IFU mosaic of the inner 0.5 pc (~125 h) maps abundance anomalies (Li enhancement, refractory depletion) that would discriminate between ordinary tidal-disruption debris chemistry (Hills 1988; Gezari 2021) and the engineered-feeding residue of H_2. We do not request this time until the trigger condition is met.

3.3 Decision criteria

Null result (expected under H_0, H_1, and mature-H_2 alike): bolometric limit $\lesssim 10^{30}$ erg s^{-1} recorded — radiative output $\lesssim 10^{-12}$ of Eddington across the contested mass range — complementing the radio Bondi-efficiency bound and jointly constraining any radiatively inefficient flow via the fundamental plane (Merloni et al. 2003; Mahida et al. 2026). A static detection consistent with a quiescent accretion SED supports H_0 and immediately sharpens every dynamical program (the source position becomes the kinematic centre). A variable detection triggers multi-messenger follow-up. A >3σ extended mid-IR excess with no stellar counterpart is the single strongest electromagnetic anomaly the campaign can produce and would be adjudicated under Section 11 rules.

4. Program 2: Radio — deep imaging, narrowband SETI, and pulsar timing

4.1 Science case

The radio program carries three loads at once. (i) *Continuum*: a deeper interferometric limit (or first detection) on the central compact source, extending Mahida et al. (2026) and the globular-cluster IMBH radio campaigns (Strader et al. 2012; Tremou et al. 2018). (ii) *Narrowband SETI*: the first dedicated technosignature search of any kind at ω Cen, closing the gap left by the FAST pilot survey's latitude limit (Huang et al. 2026). (iii) *Pulsar timing*: the binding constraint on the central mass comes from the cluster's MSP population — five discovered with Parkes (Dai et al. 2020), timed to $\dot{\nu}$ within 3.5 yr (Dai et al. 2023), then thirteen more with MeerKAT/TRAPUM (Chen et al. 2023; Stappers & Kramer 2016) for eighteen total — and the path to resolving the mass tension runs directly through more pulsars and longer baselines (Bañares-Hernández et al. 2025).

4.2 Implementation: MeerKAT era (2026–2030)

Timing. Bi-weekly L-band sessions (~26 per year, ~3–5 h each) on all eighteen MSPs with MeerKAT (Jonas 2016), producing times of arrival at ~1 μs precision for the brightest objects. Line-of-sight acceleration precision from timing scales steeply with baseline ($\sigma_a \propto T^{-5/2}$ for white noise); by analogy with the mature 47 Tuc and Terzan 5 programs (Freire et al. 2017; Prager et al. 2017), five years of MeerKAT data yield per-pulsar acceleration uncertainties of order 10^{-11}–10^{-10} m s^{-2} for the best-timed objects. For scale, a 4×10^4 $M_\odot$ point mass imposes a $\approx 6\times10^{-8}$ m s^{-2} at r = 0.3 pc: the

discriminating signal is not the detection of acceleration (already achieved; Dai et al. 2023; Bañares-Hernández et al. 2025) but the *radial profile* of accelerations and jerks across the MSP population, which distinguishes a point mass ($a \propto r^{-2}$) from an extended remnant component. Commensal search observations are expected to add 5–10 MSPs (the luminosity function is far from exhausted; Chen et al. 2023), and the point-mass upper bound tightens approximately as $1/\sqrt{N_{psr}}\,(T_0/T)^{\geq 1}$: ten well-timed pulsars with ten-year baselines push a 6,000 $M_\odot$ bound to the 3,000–4,000 $M_\odot$ level if no point-mass signal emerges — or find one.

Narrowband SETI. Commensally with every timing session, a 1-Hz-resolution spectrometer backend records the full primary beam (~1°, comfortably containing the whole cluster). Processing follows standard drift-search practice (Enriquez et al. 2017): Doppler drifts ±4 Hz s^{-1}, SNR >15 candidate threshold, on–off cadence against reference pointings, and re-detection in an independent session required before any candidate advances. Seventy hours over two years reach a minimum detectable EIRP of ~(3–6)×10^{17} W at 5.49 kpc — roughly an order of magnitude above the FAST pilot's ~10^{16} W thresholds at its (closer, northern) targets (Huang et al. 2026), but the first limit of any kind for this cluster, and sufficient to detect any transmitter exceeding ~10^4 times the EIRP of the Arecibo planetary radar — a modest output for an energy-rich civilization — from any of ~10^7 stars in the beam simultaneously.

Continuum. The summed continuum visibilities from the timing campaign provide a ~100+ h synthesis image at L/S band; combined with the ATCA 7.25-GHz limit (Mahida et al. 2026) this constrains both the flat-spectrum (jet) and steep-spectrum (pulsar) interpretations of any future central source candidate.

4.3 Implementation: SKA era (2029–)

SKA-Mid science verification is expected in 2029 and early operational cycles in the early 2030s (Braun et al. 2019). The ω Cen program transfers wholesale: timing precision improves by the sensitivity ratio (factor ~4–5 over MeerKAT for these declinations), the MSP census plausibly doubles, and the narrowband EIRP threshold drops toward ~10^{16} W — FAST-class sensitivity on a southern target FAST cannot see. By LISA launch, the pulsar acceleration map will be the best non-GW constraint on the central potential in any globular cluster.

4.4 Decision criteria

The timing program is the campaign's primary near-term discriminator between H_0 and H_1: a smooth r^{-2} acceleration profile centred on the kinematic centre at $\geq 10^4\ M_\odot$ retires H_1; a persistently tightening upper bound below the kinematic lower bounds forces revision of the fast-star analysis (membership, foreground, or binarity systematics) and shifts the campaign's centre of gravity to H_1. For SETI: any re-detected narrowband candidate enters the two-messenger protocol; a null is published as the first ω Cen technosignature limit. Paper A dependence (labelled): under H_2, P3 of that paper predicts *no* leakage radiation, so SETI nulls carry no evidential weight against H_2 — they constrain only conventional beacon scenarios.

5. Program 3: Astrometry — HST, Gaia DR4, Roman, and ELT/MICADO

5.1 The honest calculation first

The intuitively obvious astrometric test — watch the seven fast stars curve — does not work on a useful timescale, and we show this rather than assert it. The proper-motion acceleration induced by a central mass $M_\bullet$ at projected radius r is

$$a = GM_\bullet/r^2 \approx 4.4\times10^{-7}\ (M_\bullet/2\times10^4\ M_\odot)\ (r/0.08\ \mathrm{pc})^{-2}\ \mathrm{m\ s^{-2}} \quad (1)$$

which at d = 5.49 kpc corresponds to $\approx 5.3\times10^{-4}$ mas yr^{-2} (mean projection factors of order unity absorbed). Against this, the oMEGACat astrometry — ~0.02 mas single-epoch precision over a ~15-year HST baseline — delivers acceleration uncertainties of $\sigma_a \approx 1\times10^{-3}$ mas yr^{-2} per star (Häberle et al. 2024a, 2024b): a 0.5σ measurement. Extending HST monitoring to 2028 (a 26-year baseline) improves this only to ~0.7σ; a four-year ELT/MICADO campaign at ~50–100 μas per epoch (Davies et al. 2021) reaches $\sigma_a \approx 2\times10^{-3}$ mas yr^{-2} — worse, because acceleration precision scales as T^{-2} and four years is short (Figure 2). Direct 5σ curvature detection at nominal parameters requires either $M_\bullet \gtrsim 10^5\ M_\odot$, a star at r < 0.015 pc, or ≳30-year baselines on ELT-class astrometry. We therefore *retain* the fast-star monitoring (cheap; protects against the lucky cases, since González Prieto et al. 2025-mass holes and undiscovered inner stars are both live possibilities) but *route the decision-grade astrometry through three measurements that do work*, below.

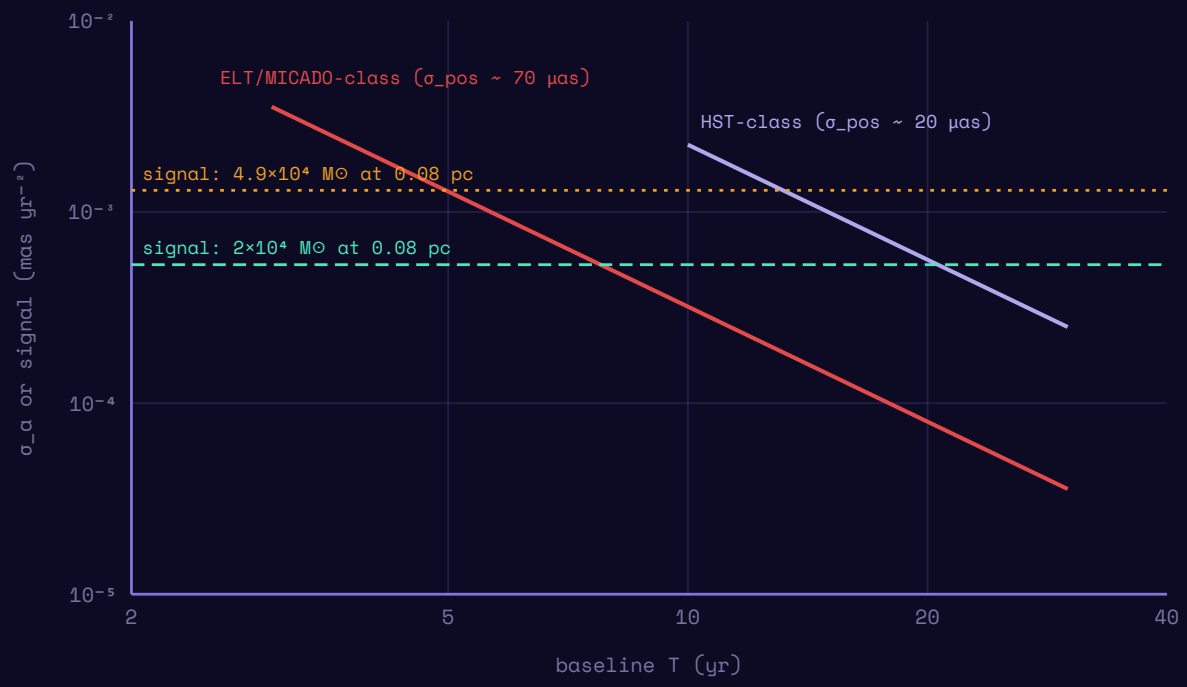

Figure 2. Why direct acceleration detection on the known fast stars is marginal: 1σ acceleration uncertainty versus baseline for HST-class and ELT-class astrometric campaigns (curves, $\sigma_a \propto \sigma_{pos} T^{-2}$, anchored to the measured oMEGACat performance and nominal MICADO performance), against the expected signals (horizontal lines, Eq. 1). Curves assume uniform observing cadence over the full baseline and are therefore optimistic for sparse epoch-limited extensions (the realistic 2028 HST extension yields only ~0.7σ). Even the heavy González Prieto et al. (2025) mass crosses 5σ only for baselines approaching two decades. The astrometric program is therefore built around wander, reference-frame, and discovery measurements instead (Section 5.2).

5.2 Three astrometric measurements that do work

(i) Photocentric wander (Roman). An IMBH of mass $M_\bullet$ in dynamical equilibrium with stars of mean mass $\bar{m}$ executes Brownian motion with velocity $\sigma_\bullet \sim \sigma_\star \sqrt{(\bar{m}/M_\bullet)}$ relative to the cluster centre of mass. The lighter the hole, the larger its wander: the predicted amplitude differs by ~17 per cent between 6,000 and 8,200 $M_\odot$ (the square root of the mass ratio), and by a factor ~2.4 between 8,200 and 4.9×10^4 $M_\odot$. Resolving it requires measuring the reflex motion of the innermost stellar distribution at the few-μas yr^{-1} level against the bulk cluster frame — precisely the regime of Roman's wide-field astrometry, which delivers ~10 μas-class differential astrometry over fields vastly larger than HST's, tying the inner arcseconds to thousands of cluster reference stars at once (Sanderson et al. 2019). Roman launches 2026 August 30; a five-year cadenced guest-observer program (ω Cen lies outside the core community surveys, so epochs must be requested, but the per-

epoch cost is small) discriminates the light-IMBH from heavy-IMBH wander regimes at the 3–5σ level, independently of individual-star orbits.

(ii) The absolute reference frame (Gaia DR4, 2026 December). The dominant systematic in all current inner-field astrometry is the tie between HST's relative frame and the absolute (ICRS) frame. Gaia DR4's 5.5-year solutions for the r > 10" cluster members (Gaia Collaboration 2016) improve proper motions by a factor ~4.5 over DR2-era ties, propagating directly into the fast-star velocity vectors — and thus into the 8,200 $M_\odot$ lower bound itself, which rests on those velocities exceeding the escape speed. This is the cheapest decisive measurement in the entire campaign: a re-derivation of the Häberle et al. (2024a) bound on the DR4 frame either firms the tension into a >5σ contradiction with the timing bound or dissolves it. Deliverable in 2027 from archival data alone.

(iii) Inner-star discovery (ELT/MICADO, 2029+). MICADO's ~10 mas resolution and 39-m aperture (Davies et al. 2021) penetrate the central arcsecond at magnitudes HST cannot reach in crowding, with the realistic prize being *new* fast stars at r < 0.02 pc. A single star at 0.01 pc raises the acceleration signal of Eq. (1) by a factor of 64, converting Figure 2's verdict from "decades" to "a few epochs" — this is how the Galactic-Centre program succeeded (GRAVITY Collaboration 2018), and ω Cen is its natural second act at $10^2\times$ lower mass. Two epochs per year from 2029; first acceleration-grade results plausible by the early 2030s.

5.3 Decision criteria

DR4 re-derivation (2027): if the fast-star velocity excess survives at >5σ, H_1 requires the MSP analysis to be systematically biased — a specific, checkable claim (radial distribution priors; Bañares-Hernández et al. 2025). Roman wander (2031): light/heavy discrimination feeds the LISA prior. MICADO inner stars (2030s): any star with measured Keplerian curvature yields $M_\bullet$ to tens of per cent, independent of statistical modelling, before LISA flies.

6. Program 4: The LISA forecast

6.1 Science case

LISA is the campaign's endgame: the only instrument that can measure the central object's mass and spin to high precision regardless of its electromagnetic state (Colpi et al. 2024; Amaro-Seoane et al.

2017). For a compact object inspiralling into a 10^4-class IMBH, matched-filter parameter estimation delivers fractional precisions of order 10^{-3}–10^{-4} on mass and 10^{-3} on spin (Babak et al. 2017; Amaro-Seoane 2018). At 5.49 kpc — versus the Gpc distances over which LISA expects to detect its EMRI population — any ω Cen inspiral in band during the mission would be loud (signal-to-noise in the hundreds to thousands), making the limiting factor *occurrence probability, not sensitivity*.

6.2 Rates, honestly

Cluster-dynamics simulations imply IMRI formation rates per cluster of order 10^{-8}–10^{-6} yr^{-1} depending on IMBH mass, binary fraction, and core density; over a 4–10-year mission this yields an in-band inspiral probability for ω Cen of order 10^{-2} or below — possibly enhanced by the cluster's unusually dense remnant population, but not to order unity. The forecast must therefore be stated in three tiers. **Tier 1 (probable):** no ω Cen source; LISA still constrains the stochastic foreground from the core remnant population, directly testing H_1's extended-component density profile. **Tier 2 (plausible):** a quasi-monochromatic source — a compact object orbiting the IMBH well outside plunge — detectable over the mission as a continuous-wave source; this measures $M_\bullet$ through the orbital frequency and its drift, to precision intermediate between timing and a full inspiral. **Tier 3 (lucky):** a genuine inspiral; per-mille mass and spin. The campaign's structure ensures that even Tier 1 is decisive when combined with the pulsar acceleration map (Section 4): a stochastic-foreground non-detection plus a point-mass timing profile isolates H_0.

6.3 The spin measurement as the H_2 adjudicator

Paper A dependence (labelled): under H_2 the single sharpest residue is near-extremal spin ($a_\star \gtrsim 0.9$) on a hole that has demonstrably accreted nothing for Gyr; under H_0 natural formation channels for cluster IMBHs span low-to-moderate spin. A Tier 3 spin measurement therefore adjudicates H_2 at the same instant it completes the H_0/H_1 question — the entire reason this campaign's electromagnetic and dynamical groundwork must precede LISA rather than follow it. The full decision tree, including this branch, is assembled in Section 13.

6.4 Deliverables before launch

Between now and the mid-2030s the program's work is preparatory but concrete: (i) maintain the dynamical model of the inner parsec (timing + astrometry) so that any LISA source has an immediate

host context; (ii) publish the ω Cen-specific EMRI/IMRI rate forecast with the post-2026 mass constraints folded in; (iii) ensure the cluster's barycentric ephemeris and distance (±0.06 kpc; Häberle et al. 2025) are maintained at the precision LISA parameter estimation will assume.

7. Program 5: KM3NeT/ARCA neutrino monitoring

7.1 Science case and geometry

ω Cen at declination −47.5° is an *up-going* source for the Mediterranean: KM3NeT/ARCA observes it through the Earth, with atmospheric muons filtered out and sub-0.2° angular resolution for track events at ≳10 TeV (Adrián-Martínez et al. 2016). For IceCube the same source is down-going (degraded but usable above ~100 TeV; IceCube Collaboration 2020); the combined ANTARES+IceCube southern-sky analysis (Albert et al. 2020) defines the archival baseline. ARCA stood at 51 of 230 detection units in early 2026, with real-time alert distribution entering commissioning; sensitivity grows roughly linearly with instrumented volume through the decade. The conventional science driver is a deep point-source limit on a dense old stellar system; the technosignature driver (Paper A dependence, labelled) is the burst-mode prediction of black-hole-based computation models (Dvali & Osmanov 2023; Lacki & DiKerby 2025), for which a cluster like ω Cen is the natural test bed and which the existing point-source pipelines do not target: their time-integrated statistics dilute rare short multiplets.

7.2 Implementation

Steady-state search (archival + ongoing). Standard unbinned-likelihood point-source analysis (Braun et al. 2008) at the ω Cen coordinates in (i) the ANTARES+IceCube combined dataset, (ii) the growing ARCA exposure. Full-array ARCA reaches $E^2\Phi$ ~ few $\times 10^{-12}$–10^{-11} TeV cm^{-2} s^{-1} at this declination over multi-year integrations — the deepest neutrino limit ever placed on a globular cluster.

Burst pipeline (new, the campaign's contribution). A pre-registered transient search in sliding windows of 10^2, 10^3, and 10^4 s within 1° of the cluster centre, E ≳ 10 TeV, on ARCA data plus IceCube alerts. The detection criterion — ≥3 tracks in one window, ≥5σ post-trials, no plausible astrophysical counterpart — is calibrated by the background expectation: the atmospheric-neutrino rate above 10

TeV within a 1° cone is of order 10^{-7} s^{-1}, so a 10^3-s window expects ~10^{-4} events and a triplet is intrinsically far beyond 5σ *before* trials; the design problem is the trials budget over a decade of monitoring (~3×10^6 windows), which the post-trials threshold absorbs (Figure 3). Any KM3NeT or IceCube real-time alert within 1° triggers the electromagnetic protocol of Section 11 — the alert system's first ω Cen-relevant cycle begins in late 2026.

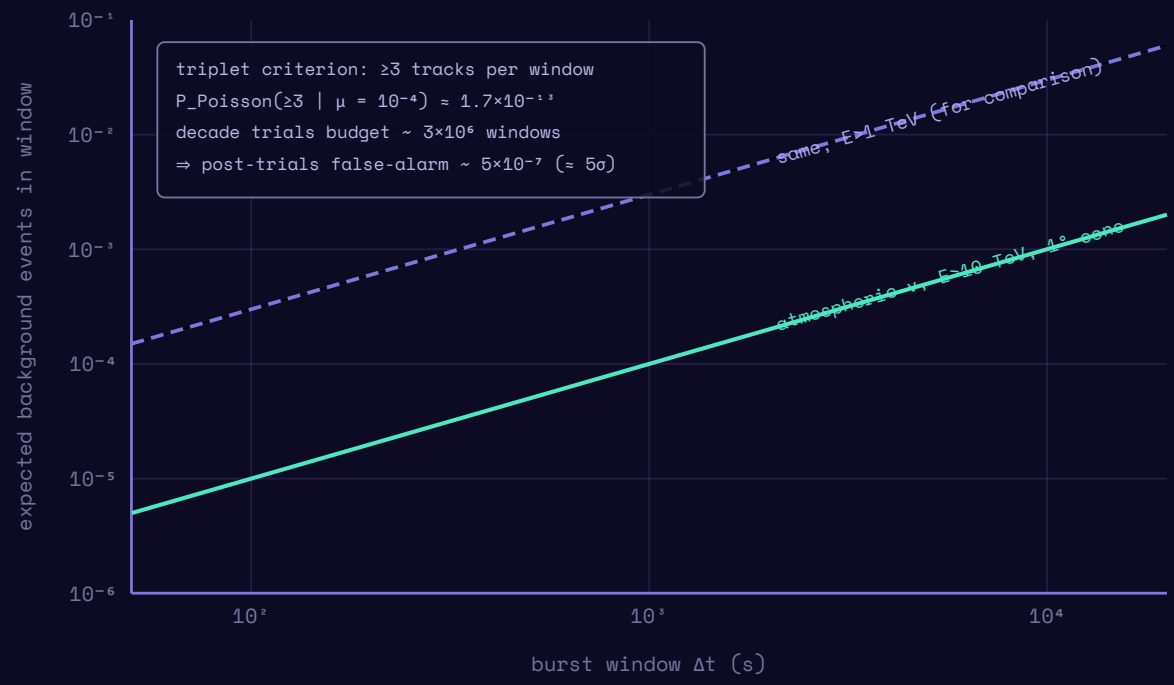

Figure 3. Design of the neutrino burst search. Expected atmospheric-neutrino background within 1° of ω Cen versus window length, for the working E > 10 TeV threshold (solid) and a softer 1 TeV threshold (dashed). At 10^3 s the expectation is ~10^{-4} events, so the pre-registered triplet criterion is intrinsically secure against background even after a decade-scale trials budget; rate normalizations are order-of-magnitude, with exact values set by the detector Monte Carlo at analysis time.

7.3 Decision criteria

A decade-scale null at full-array sensitivity closes the burst-mode channel of Dvali & Osmanov (2023) for this system at its predicted strength — a publishable falsification regardless of H_2 (the underlying micro-black-hole physics is itself contested; Álvarez-Domínguez et al. 2024; Loeb 2024). One confirmed multiplet with no counterpart, repeated at the same location, would be among the most significant anomalies in high-energy astrophysics independent of any ETI interpretation; the protocol deliberately treats it as an astrophysical discovery first.

8. Program 6: Gamma rays — Fermi-LAT archival and CTAO-South

8.1 Science case

ω Cen is a GeV source: 4FGL J1326.7−4729 coincides with the core, and its emission is attributed to the MSP ensemble after deep searches found no individual pulsations (Dai et al. 2020, 2023). This is both foreground and opportunity: the steady MSP emission is a calibration source against which *transient* or *spectrally anomalous* excesses can be sought with high contrast. The technosignature channel shares the burst phenomenology of Program 5 (the Dvali & Osmanov 2023 spectrum is "democratic" across species, implying γ-ray counterparts to any neutrino burst); the conventional channel is a TeV detection or limit on the cluster, of independent interest for MSP magnetospheric and inverse-Compton physics.

8.2 Implementation

Fermi-LAT (archival, immediate). Re-analysis of the now >17-year LAT exposure at the cluster position: (i) refit of the steady source with current diffuse models; (ii) a sliding-window transient search ($\Delta t = 10^2$–10^4 s, matched to Program 5's windows) over the full mission, with the steady MSP flux as null model; (iii) spectral-cutoff analysis — MSP ensembles cut off at a few GeV, so any significant emission above ~30 GeV flags a non-MSP component. All three are pure archival analyses.

CTAO-South (2027+). The first Paranal telescopes arrive through late 2026 (CTA Consortium 2019); intermediate arrays operate while construction proceeds. We propose (i) a 50-h deep exposure reaching ~10^{-13} erg cm^{-2} s^{-1} at 1 TeV — the first meaningful TeV limit on any globular cluster core at this depth — and (ii) a standing target-of-opportunity trigger: any Program-5 alert within 1° pre-authorizes prompt CTAO follow-up, the latency being minutes for a source that transits Paranal ~23° from zenith.

8.3 Decision criteria

A LAT transient coincident with a neutrino window elevates any Program-5 candidate to the campaign's highest tier. A TeV detection of steady emission is conventional astrophysics (and would be the first for a globular cluster); its absence at CTAO depth tightens the inverse-Compton budget of

the MSP population. Spectral excess above 30 GeV without transient behaviour prompts a dark-matter-versus-exotica analysis before any H_2 language is entertained.

9. Program 7: Optical/IR time domain — Rubin/LSST and archival anomalies

9.1 Science case

Two distinct searches share this program. First, *transients*: tidal disruption of a star by an IMBH is the one natural event that would announce the central object unambiguously (Hills 1988; Gezari 2021); rates are low (~10^{-7}–10^{-6} yr^{-1} per cluster for full disruptions, higher for partial-disruption flares of the dense remnant environment) but the cost of monitoring is now effectively zero, because LSST began surveying in early 2026 and ω Cen sits in the survey footprint (Ivezić et al. 2019). Second, *archival photometric anomalies*: the oMEGACat HST archive — 500+ epochs over 20+ years for 10^6 stars (Häberle et al. 2024b; Nitschai et al. 2023) — is an essentially unmined time-domain dataset in which periodic dimmings, secular fades, or statistically anomalous variability classes (the photometric residues conjectured for megastructures; Wright et al. 2014; Hsiao et al. 2021; Socas-Navarro et al. 2021) can be sought with two decades of baseline for free.

9.2 Implementation

LSST stream. The cluster core saturates LSST, but the half-light region and outskirts do not: a standing broker filter on the ω Cen field flags (i) any nuclear transient brighter than r ~ 16 (TDE flares at 5.49 kpc would reach naked-eye-adjacent magnitudes; even heavily extincted or partial events are unmissable), and (ii) anomalous variables among ~10^5 measurable members. Difference-imaging in the crowded annulus is handled with the oMEGACat reference catalogue.

Archival mining. A systematic variability census of the oMEGACat photometric archive: period search to ~0.05 mag depth across 10^6 stars, secular-trend extraction, and an explicitly pre-registered anomaly statistic (population-level excess of non-physical light-curve classes) with human vetting of survivors. Conventional yield: the deepest variable-star census of any globular cluster, eclipsing-binary distance cross-checks, and stellar-rotation demographics — publishable regardless of anomalies.

Core-depletion statistic (joint with Program 3). Paper A dependence (labelled): H_2's P4c predicts secular depletion of the loosely bound core population beyond N-body expectations. The measurement — star counts by binding-energy quantile against matched N-body models (González Prieto et al. 2025) — is identical to standard mass-segregation analysis and is conventional dynamics on its face; we simply pre-register the depletion direction as a tracked statistic.

9.3 Decision criteria

A nuclear TDE settles the IMBH question electromagnetically (and its decay light curve constrains $M_\bullet$; Gezari 2021); the campaign's role is to guarantee the multi-wavelength response is pre-arranged. Archival anomalies are population statistics: no single weird star means anything (the false-positive history of photometric SETI is long); only a $\geq 3\sigma$ class-level excess, surviving vetting, enters the coordination protocol.

10. Program 8: Supplementary archival channels

Two channels are cheap enough to run as graduate-student archival projects and complete the messenger coverage.

Continuous gravitational waves (LIGO–Virgo–KAGRA). Narrowband searches at twice the spin frequencies of the eighteen known MSPs, using public O4/O5 data and the MeerKAT ephemerides, following the known-pulsar methodology of Abbott et al. (2022). Expected upper limits $h_0 \lesssim 10^{-26}$ translate to ellipticity limits $\varepsilon \lesssim$ few $\times 10^{-8}$ at 5.49 kpc for the faster spinners — conventional neutron-star physics in an environment (a dense old cluster) where the recycled-pulsar population is dynamically distinctive. The technosignature reading (structured masses in MSP orbits) is not separately funded; it shares the identical data product.

Ultra-high-energy cosmic rays (Pierre Auger). ω Cen sits well inside Auger's southern acceptance. A stacked directional analysis around the cluster position in the public $E > 8$ EeV dataset (Pierre Auger Collaboration 2017), with 1° and 3° apertures bracketing magnetic deflection, either contributes an upper limit on hadronic acceleration in the core — of conventional interest given the MSP wind population — or flags an anisotropy for which the Galactic-Centre region (18° away) is the obvious confounder to exclude first.

11. Coordination, joint statistics, and data policy

11.1 The two-messenger rule, operationalized

Table 3 pre-registers the campaign's trigger and confirmation matrix. Each row names a primary trigger, the channels that can confirm it, and the joint-significance requirement. Combined significances across independent channels use Fisher's method on the per-channel p-values, with the global claim threshold set at $p < 5\times10^{-7}$ post-trials — deliberately at discovery convention, because the base rate of true technosignatures is unknown and plausibly zero, so the loss function is asymmetric: a false positive damages the field's credibility far more than a delayed true positive damages the discovery (Wright et al. 2022).

Table 3. Pre-registered trigger and confirmation matrix. "Joint" = Fisher-combined post-trials significance required for any anomaly claim; single-channel events, however significant, are reported as astrophysical candidates only.

PRIMARY TRIGGER	CHANNEL THRESHOLD	CONFIRMING CHANNELS	DISPOSITION
Neutrino multiplet (P5)	≥3 tracks / 1° / $10^{2–3}$ s / ≥10 TeV, 5σ post-trials	Fermi-LAT window (P6); CTAO ToO (P6); JWST/ground IR (P1)	repeat at same position required for anomaly claim; single event published as astrophysical transient
Narrowband radio line (P2)	SNR >15, re-detected in independent session	off-source cadence; second telescope (e.g. ATCA)	interference exclusion protocol precedes any claim
IR variable/transient (P1, P7)	≥20% amplitude or new source >5σ	radio continuum (P2); X-ray archival; LSST stream (P7)	classified as accretion/TDE candidate first
Photometric anomaly class (P7)	≥3σ population-level excess	independent archive (Gaia epochs); spectroscopy	population statistic only; no single-star claims
Astrometric anomaly (P3)	core-depletion or wander outside H_0/H_1 envelopes at 3σ	timing acceleration map (P2)	feeds LISA prior; no standalone claim

PRIMARY TRIGGER	CHANNEL THRESHOLD	CONFIRMING CHANNELS	DISPOSITION
LISA spin/mass result (P4)	per mission parameter estimation	all of the above as context	adjudicates $H_0/H_1/H_2$ per Section 13

11.2 Alert latency and standing authorizations

The latency-critical path is neutrino → gamma/optical: KM3NeT real-time alerts (commissioning through 2026) distribute within seconds–minutes; the CTAO ToO and LSST-broker hooks are standing authorizations requiring no human decision for the first response hour. Slow-path coordination (JWST DDT requests, radio re-pointing) is pre-drafted as template proposals with trigger criteria attached, cutting submission latency to <24 h.

11.3 Data policy

All pipelines, thresholds, and analysis code are public from the start (the pre-registration is itself the methods paper); all derived catalogues are released with their papers; raw data inherit host-facility policies. Negative results are published on a fixed cadence — the field's credibility problem is survivorship bias, and a campaign explicitly designed around falsification should model the cure.

11.4 Program summary

Table 4 assembles the eight programs with sensitivities, requests, and timelines.

Table 4. Campaign summary. "Type": A = archival/commensal, N = new observations, F = free by-product of funded survey. Costs are rough full-cost estimates (Section 12); facility time is listed separately since it is allocated, not purchased.

#	PROGRAM	TYPE	KEY SENSITIVITY	FACILITY REQUEST	TIMELINE	DECIDES
P1	JWST IR limits	N	~10 nJy (F444W); $L_{bol} \lesssim 10^{30}$ erg s^{-1}	~65 h, 2 epochs	2027–2029	accretion state; waste heat
P2	Radio: timing + SETI + continuum	N	$\sigma_a \sim 10^{-11}$–10^{-10} m s^{-2} (5 yr); EIRP (3–6)×10^{17} W	~130 h yr^{-1} MeerKAT, 5 yr	2026–2031+	H_0 vs H_1 (profile); first SETI limit

#	PROGRAM	TYPE	KEY SENSITIVITY	FACILITY REQUEST	TIMELINE	DECIDES
P3	Astrometry: DR4 + Roman + ELT	A/F/N	frame to DR4; wander 3–5σ (light vs heavy); inner-star discovery	archival; Roman survey; ~28 h ELT	2026–2035	mass tension; wander mass
P4	LISA forecast	F	$\delta M/M \sim 10^{-3-4}$, $\delta a_\star \sim 10^{-3}$ (if inspiral)	none (mission science)	2035+	everything (Tier 3)
P5	Neutrino monitoring	A/N	$E^2\Phi \sim$ few$\times 10^{-12}$ TeV cm^{-2} s^{-1}; burst triplets	pipeline on ARCA stream	2026–2040	Dvali–Osmanov channel
P6	Gamma rays	A/N	LAT transient windows; 10^{-13} erg cm^{-2} s^{-1} at 1 TeV (50 h)	50 h CTAO + ToO	2026–2032	burst counterparts; TeV first
P7	Time domain	A/F	TDE flares unmissable; 10^6-star archival census	LSST stream; archival	2026–2036	TDE; anomaly classes
P8	CW + UHECR archival	A	$h_0 \lesssim 10^{-26}$; stacked UHECR limit	public data	2026–2029	completeness

12. Cost and timeline

12.1 Costing

Table 5 presents full-cost estimates (salaries, computing, travel, publication; US academic rates) for each program. Facility time is excluded — it is competitively allocated and carries no cash cost to the campaign — as is LISA, which is mission science. The total, $\lesssim$ US$7M over a decade, is calibrated against the per-program budgets developed in the precursor proposal studies for this campaign and is small by every relevant comparison: a single mid-scale NASA Explorer instrument, one JWST cycle's archival-funding pool, or the marginal cost of a few nights of ELT operations.

Table 5. Indicative full-cost budget (US$k, 2026 dollars). Personnel figures assume postdoc-led programs with PI/Co-I fractions and graduate students where noted. Estimates for the Roman component of P3 and for P5, P6, P7 are new to this paper; others follow the precursor proposal studies.

PROGRAM	COST	YEARS	DOMINANT ELEMENT
P1 JWST deep imaging + waste heat	1,310	3	postdoc + grad + HPC
P1b JWST spectroscopic phase (contingent)	257	2	postdoc (triggered only)
P2 MeerKAT timing + SETI + continuum	2,900	5	2 postdocs + grad + computing
P3 HST/Gaia frame + Roman wander	425	4	postdoc + HPC
P3b ELT/MICADO epochs	490	4	postdoc share + travel
P4 LISA preparatory modelling	150	3	PI/Co-I fractions
P5 Neutrino burst pipeline	250	3	postdoc + compute
P6 Fermi archival + CTAO ToO	200	3	postdoc share
P7 LSST broker + archival mining	300	3	postdoc + broker engineering
P8 CW + UHECR archival	239	3	graduate RA + HPC
Coordination, data releases, workshops	150	10	part-time coordinator
Total (P1b excluded)	**6,414**	—	—

12.2 Timeline

Figure 4 lays the programs against the facility schedule. The structure is deliberate: the 2026–2028 window is dominated by archival work and standing-pipeline construction (cheap, immediate); 2029–2034 by the new-facility harvest (ELT, SKA-Mid, CTAO completion, Roman mid-mission); 2035+ by LISA. Decision points D1–D4 are marked and defined in Section 13.

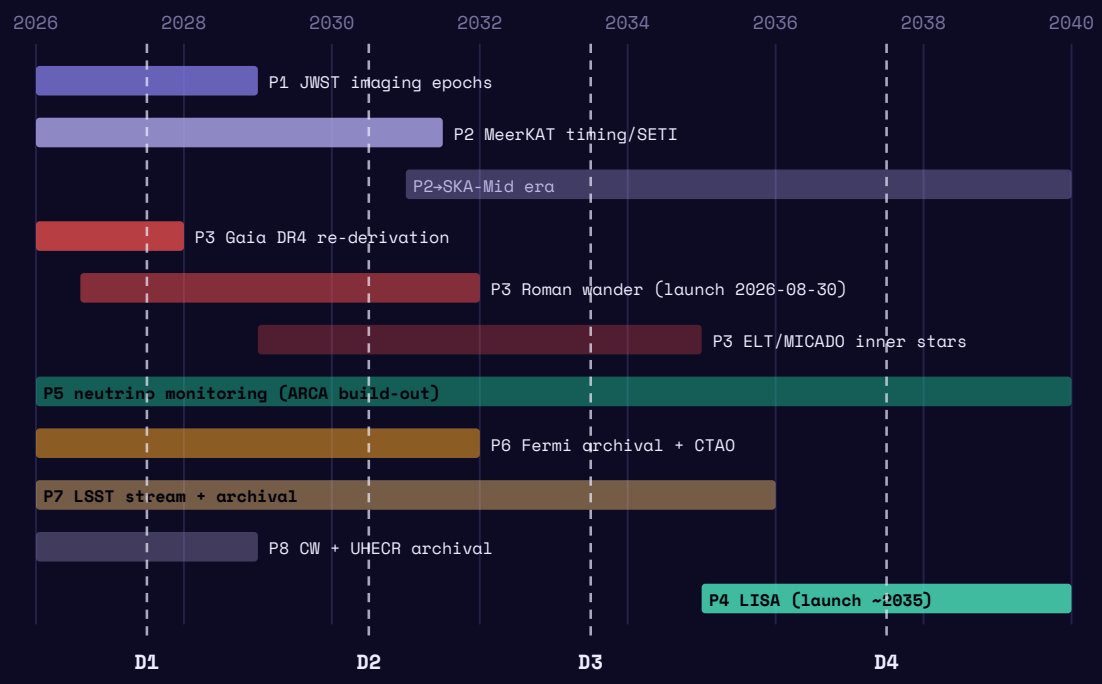

Figure 4. Campaign timeline against the facility schedule, 2026–2040. Decision points: D1 (2027): Gaia-DR4 re-derivation of the fast-star bound. D2 (2030): five-year pulsar acceleration profile; Roman wander first results. D3 (2033): pre-LISA synthesis — H_0 versus H_1 called at available significance. D4 (2037): LISA verdict, if a source is caught.

13. The decision structure through 2040

Figure 5 assembles the campaign into a single decision tree, extending the gravitational-wave branch of Paper A's falsification framework backward into the electromagnetic-and-timing era. Rough branch weights are stated where the inputs exist to estimate them; they are planning priors, not forecasts, and we expect D1/D2 to revise them substantially.

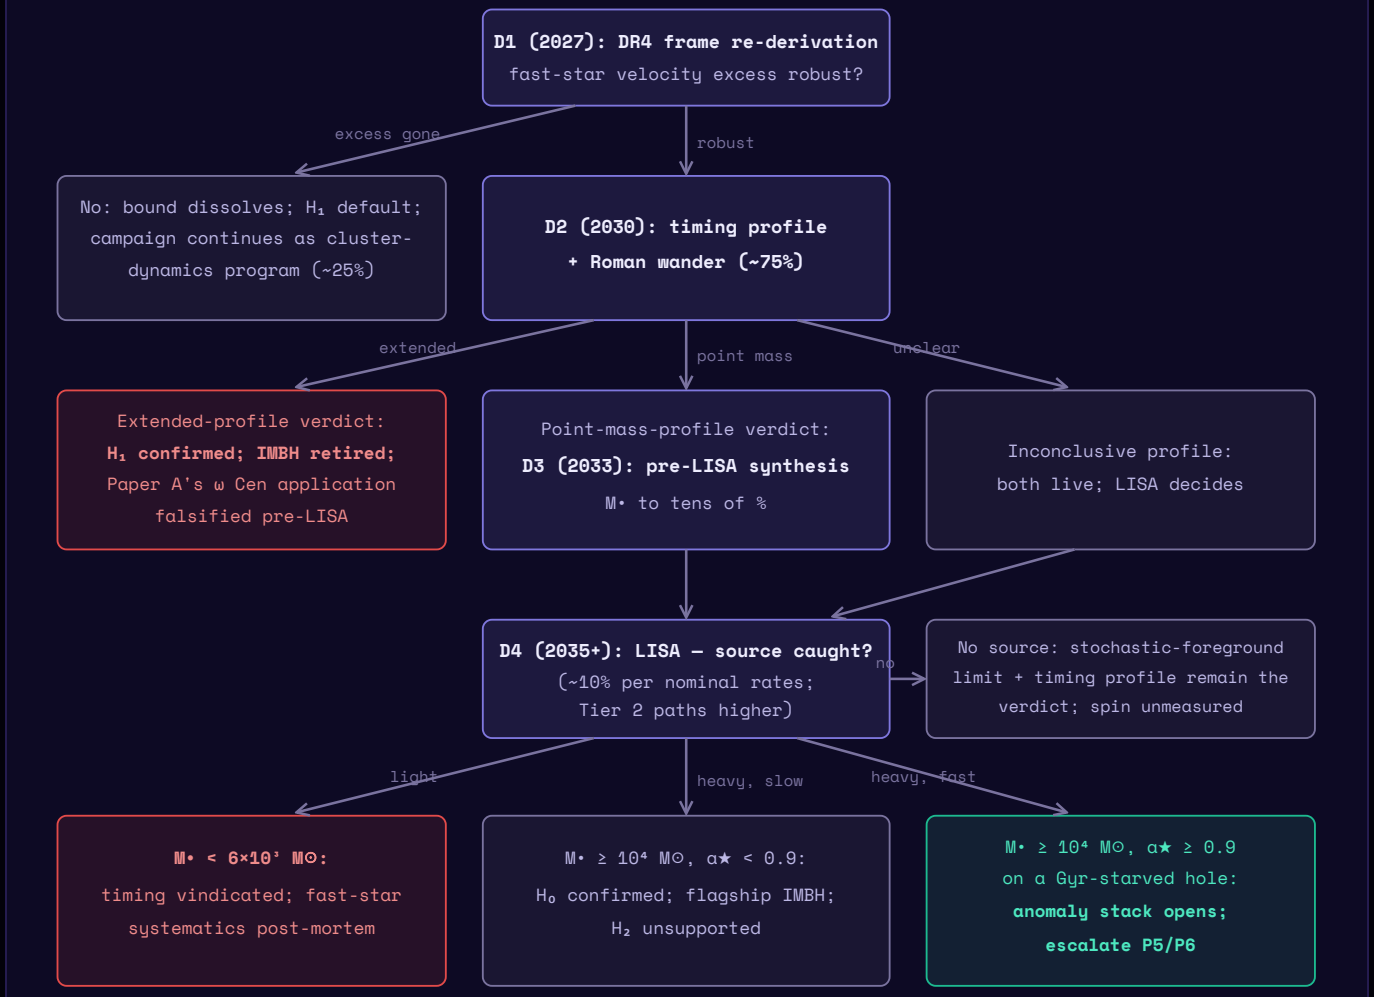

Figure 5. Decision structure through 2040. Red boxes retire hypotheses; the single green branch is the only path on which the technosignature interpretation gains material support — and even there the claim is "anomaly requiring explanation." Percentages are planning priors only. The tree's essential property is that *every* terminal node is a publishable scientific result about a first-rank astrophysical object.

14. Conclusion

Omega Centauri in 2026 presents a configuration that observational astronomy rarely supplies: a contested discovery (the Galaxy's best IMBH candidate) whose two strongest constraints formally contradict each other; an electromagnetic silence deep enough to be a result in itself; an essentially untouched technosignature parameter space; and a fleet of new southern and all-sky instruments — Rubin, Roman, Gaia DR4, KM3NeT, CTAO, SKA-Mid, ELT, and ultimately LISA — arriving on exactly the timescale needed to resolve all of it. The campaign presented here is, at bottom, an exercise in not wasting that conjunction.

Its architecture reflects three commitments. First, *dual use*: every program stands as conventional astrophysics — the mass tension, the quiescent-accretion frontier, the MSP dynamical laboratory, the first deep TeV and neutrino limits on a globular cluster — so that no null result is a loss. Second, *honesty about sensitivity*: where the obvious measurement fails (direct fast-star acceleration before ~2040), we demonstrate the failure quantitatively and redesign around it, because a program that oversells its tests invites the credibility collapse that has repeatedly damaged this field. Third, *pre-registration*: hypotheses $H_0/H_1/H_2$, thresholds, confirmation rules, and decision points are fixed in advance, so that whatever the sky delivers — a remnant swarm, a quiet heavy hole, or a genuine anomaly — the inference chain was written down before anyone knew the answer.

The cost is under US$7M across a decade, most of it salaries for archival analysis riding on facilities that are funded regardless. The payoff structure is asymmetric in the campaign's favour: the probable outcomes (a resolved mass tension, a characterized central object, the first technosignature limits on 10^7 old stars) are solid science at modest cost, and the improbable outcome — the green branch of Figure 5 — would justify the program many thousands of times over. We commit, symmetrically, to the red branches: if the timing profile confirms an extended remnant component, or LISA delivers a light or slowly spinning hole, the anomaly hypotheses retire here without special pleading, and what remains is what was always underneath: the most interesting stellar system in the southern sky, finally measured properly.

Acknowledgements and disclosure

The author thanks the maintainers of the NASA Astrophysics Data System and arXiv, on which the citation verification for this work relied. **AI assistance disclosure:** drafting, citation verification, derivation checking, and figure preparation for this manuscript were performed with substantial assistance from a large language model (Claude, Anthropic), under the author's direction; the author reviewed and takes full responsibility for all claims, derivations, and references. Interactive calculators implementing the quantitative material in this paper, together with the underlying per-instrument proposal studies, are available at omegacentauri.me.

Data availability

No new observational data were generated for this work. All quantitative claims derive from the cited literature; analysis conventions for the proposed programs will be released with the respective pipeline papers.

References

All references verified against ADS/arXiv/publisher records, June 2026. Preprints are marked as such.

1. Abbott, R., et al. (LIGO–Virgo–KAGRA Collaborations) (2022). Searches for gravitational waves from known pulsars at two harmonics in the second and third LIGO-Virgo observing runs. *ApJ*, 935, 1. doi:10.3847/1538-4357/ac6acf · arXiv:2111.13106
2. Adrián-Martínez, S., et al. (KM3NeT Collaboration) (2016). Letter of intent for KM3NeT 2.0. *J. Phys. G*, 43, 084001. doi:10.1088/0954-3899/43/8/084001 · arXiv:1601.07459
3. Albert, A., et al. (ANTARES & IceCube Collaborations) (2020). ANTARES and IceCube combined search for neutrino point-like and extended sources in the southern sky. *ApJ*, 892, 92. doi:10.3847/1538-4357/ab7afb · arXiv:2001.04412
4. Álvarez-Domínguez, Á., Garay, L. J., Martín-Martínez, E., & Polo-Gómez, J. (2024). No black holes from light. *Phys. Rev. Lett.*, 133, 041401. doi:10.1103/PhysRevLett.133.041401 · arXiv:2405.02389
5. Amaro-Seoane, P., et al. (2017). Laser Interferometer Space Antenna. arXiv:1702.00786 (LISA L3 mission proposal submitted to ESA)
6. Amaro-Seoane, P. (2018). Relativistic dynamics and extreme mass ratio inspirals. *Living Rev. Relativ.*, 21, 4. doi:10.1007/s41114-018-0013-8
7. Babak, S., et al. (2017). Science with the space-based interferometer LISA. V. Extreme mass-ratio inspirals. *Phys. Rev. D*, 95, 103012. doi:10.1103/PhysRevD.95.103012 · arXiv:1703.09722
8. Bañares-Hernández, A., Calore, F., Martín Camalich, J., & Read, J. I. (2025). New constraints on the central mass contents of Omega Centauri from combined stellar kinematics and pulsar timing. *A&A*, 693, A104. doi:10.1051/0004-6361/202451763 · arXiv:2408.00939
9. Baumgardt, H., & Hilker, M. (2018). A catalogue of masses, structural parameters, and velocity dispersion profiles of 112 Milky Way globular clusters. *MNRAS*, 478, 1520–1557. doi:10.1093/mnras/sty1057 · arXiv:1804.08359
10. Baumgardt, H., & Vasiliev, E. (2021). Accurate distances to Galactic globular clusters through a combination of Gaia EDR3, HST, and literature data. *MNRAS*, 505, 5957–5977. doi:10.1093/mnras/stab1474 · arXiv:2105.09526
11. Braun, J., Dumm, J., De Palma, F., Finley, C., Karle, A., & Montaruli, T. (2008). Methods for point source analysis in high energy neutrino telescopes. *Astropart. Phys.*, 29, 299–305. doi:10.1016/j.astropartphys.2008.02.007 · arXiv:0801.1604
12. Braun, R., Bonaldi, A., Bourke, T., Keane, E., & Wagg, J. (2019). Anticipated performance of the Square Kilometre Array — Phase 1 (SKA1). doi:10.48550/arXiv.1912.12699 · arXiv:1912.12699

13. Breen, P. G., & Heggie, D. C. (2013). Dynamical evolution of black hole subsystems in idealized star clusters. *MNRAS*, 432, 2779–2797. doi:10.1093/mnras/stt628 · arXiv:1304.3401

14. Chen, S., et al. (2025). The intermediate mass black hole in Omega Centauri: constraints on accretion from JWST. arXiv:2511.20945 (preprint, submitted to ApJ)

15. Chen, W., Freire, P. C. C., Ridolfi, A., Barr, E. D., Stappers, B., Kramer, M., et al. (2023). MeerKAT discovery of 13 new pulsars in Omega Centauri. *MNRAS*, 520, 3847–3856. doi:10.1093/mnras/stad029 · arXiv:2301.03864

16. Cherenkov Telescope Array Consortium (2019). *Science with the Cherenkov Telescope Array*. World Scientific. doi:10.1142/10986 · arXiv:1709.07997

17. Clontz, C., et al. (2024). oMEGACat IV. Constraining the ages of Omega Centauri subgiant branch stars with HST and MUSE. *ApJ*, 977, 14. doi:10.3847/1538-4357/ad8621 · arXiv:2409.13855

18. Colpi, M., et al. (2024). LISA definition study report. arXiv:2402.07571 (ESA-SCI-DIR-RP-002)

19. Dai, S., et al. (2020). Discovery of millisecond pulsars in the globular cluster Omega Centauri. *ApJL*, 888, L18. arXiv:1912.08079

20. Dai, S., et al. (2023). Timing of pulsars in the globular cluster Omega Centauri. *MNRAS*, 521, 2616–2622. arXiv:2303.02834

21. Davies, R., et al. (2021). MICADO: the Multi-Adaptive Optics Camera for Deep Observations. *The Messenger*, 182, 17–21. doi:10.18727/0722-6691/5217

22. Dvali, G., & Osmanov, Z. N. (2023). Black holes as tools for quantum computing by advanced extraterrestrial civilizations. *Int. J. Astrobiology*, 22, 617–640. doi:10.1017/S1473550423000186 · arXiv:2301.09575

23. Dyson, F. J. (1960). Search for artificial stellar sources of infrared radiation. *Science*, 131, 1667–1668. doi:10.1126/science.131.3414.1667

24. Enriquez, J. E., et al. (2017). The Breakthrough Listen search for intelligent life: 1.1–1.9 GHz observations of 692 nearby stars. *ApJ*, 849, 104. doi:10.3847/1538-4357/aa8d1b · arXiv:1709.03491

25. Event Horizon Telescope Collaboration (2022). First Sagittarius A* Event Horizon Telescope results. I. *ApJL*, 930, L12. doi:10.3847/2041-8213/ac6674

26. Freire, P. C. C., Ridolfi, A., Kramer, M., Jordan, C., Manchester, R. N., Torne, P., et al. (2017). Long-term observations of the pulsars in 47 Tucanae — II. Proper motions, accelerations and jerks. *MNRAS*, 471, 857–876. doi:10.1093/mnras/stx1533 · arXiv:1706.04908

27. Gaia Collaboration (Prusti, T., et al.) (2016). The Gaia mission. *A&A*, 595, A1. doi:10.1051/0004-6361/201629272 · arXiv:1609.04153

28. Gezari, S. (2021). Tidal disruption events. *ARA&A*, 59, 21–58. doi:10.1146/annurev-astro-111720-030029 · arXiv:2104.14580

29. González Prieto, E., Rodriguez, C. L., & Cabrera, T. (2025). Growing the intermediate-mass black hole in Omega Centauri. *ApJL*, 990, L69. doi:10.3847/2041-8213/adfd4a · arXiv:2507.06316

30. GRAVITY Collaboration (2018). Detection of the gravitational redshift in the orbit of the star S2 near the Galactic centre massive black hole. *A&A*, 615, L15. doi:10.1051/0004-6361/201833718 · arXiv:1807.09409

31. Greene, J. E., Strader, J., & Ho, L. C. (2020). Intermediate-mass black holes. *ARA&A*, 58, 257–312. doi:10.1146/annurev-astro-032620-021835 · arXiv:1911.09678

32. Häberle, M., et al. (2024a). Fast-moving stars around an intermediate-mass black hole in ω Centauri. *Nature*, 631, 285–288. doi:10.1038/s41586-024-07511-z · arXiv:2405.06015

33. Häberle, M., et al. (2024b). oMEGACat II. Photometry and proper motions for 1.4 million stars in Omega Centauri and its rotation in the plane of the sky. *ApJ*, 970, 192. doi:10.3847/1538-4357/ad47f5 · arXiv:2404.03722

34. Häberle, M., et al. (2025). oMEGACat VI. Analysis of the overall kinematics of Omega Centauri in 3D: velocity dispersion, kinematic distance, anisotropy, and energy equipartition. *ApJ*, 983, 95. doi:10.3847/1538-4357/adbe67 · arXiv:2503.04903

35. Harris, W. E. (1996). A catalog of parameters for globular clusters in the Milky Way. *AJ*, 112, 1487 (2010 edition: arXiv:1012.3224). doi:10.1086/118116

36. Hilker, M., & Richtler, T. (2000). ω Centauri — a former nucleus of a dissolved dwarf galaxy? New evidence from Strömgren photometry. *A&A*, 362, 895–909. arXiv:astro-ph/0008500

37. Hills, J. G. (1988). Hyper-velocity and tidal stars from binaries disrupted by a massive Galactic black hole. *Nature*, 331, 687–689. doi:10.1038/331687a0

38. Hsiao, T. Y.-Y., et al. (2021). A Dyson sphere around a black hole. *MNRAS*, 506, 1723–1732. doi:10.1093/mnras/stab1832 · arXiv:2106.15181

39. Huang, B.-L., Tao, Z.-Z., Zhang, T.-J., & Gajjar, V. (2026). The FAST-SETI Milky Way globular cluster survey. I. *AJ*, 171, 51. doi:10.3847/1538-3881/ae2470 · arXiv:2511.21085

40. Ibata, R. A., et al. (2019). Identification of the long stellar stream of the prototypical massive globular cluster ω Centauri. *Nature Astronomy*, 3, 667–672. doi:10.1038/s41550-019-0751-x · arXiv:1902.09544

41. IceCube Collaboration (Aartsen, M. G., et al.) (2020). Time-integrated neutrino source searches with 10 years of IceCube data. *Phys. Rev. Lett.*, 124, 051103. doi:10.1103/PhysRevLett.124.051103 · arXiv:1910.08488

42. Ivezić, Ž., et al. (2019). LSST: from science drivers to reference design and anticipated data products. *ApJ*, 873, 111. doi:10.3847/1538-4357/ab042c · arXiv:0805.2366

43. Jonas, J. L., & MeerKAT Team (2016). The MeerKAT radio telescope. *PoS*, MeerKAT2016, 001. doi:10.22323/1.277.0001

44. KM3NeT Collaboration (2025). Observation of an ultra-high-energy cosmic neutrino with KM3NeT. *Nature*, 638, 376–382. doi:10.1038/s41586-024-08543-1

45. Lacki, B. C., & DiKerby, S. (2025). Possibilities for SETI at high energy. arXiv:2506.16351 (white paper, preprint)

46. Loeb, A. (2024). Comment on "No black holes from light". arXiv:2408.06714 (preprint; authors' reply at arXiv:2408.11097)

47. Mahida, A. D., et al. (2026). No evidence for accretion around the intermediate-mass black hole in Omega Centauri. *ApJ*, 996, 122. doi:10.3847/1538-4357/ae2ad4 · arXiv:2512.09649

48. Merloni, A., Heinz, S., & di Matteo, T. (2003). A fundamental plane of black hole activity. *MNRAS*, 345, 1057–1076. doi:10.1046/j.1365-2966.2003.07017.x · arXiv:astro-ph/0305261

49. Nitschai, M. S., et al. (2023). oMEGACat. I. MUSE spectroscopy of 300,000 stars within the half-light radius of ω Centauri. *ApJ*, 958, 8. doi:10.3847/1538-4357/acf5db · arXiv:2309.02503

50. Noyola, E., Gebhardt, K., & Bergmann, M. (2008). Gemini and Hubble Space Telescope evidence for an intermediate-mass black hole in ω Centauri. *ApJ*, 676, 1008–1015. doi:10.1086/529002 · arXiv:0801.2782

51. Pierre Auger Collaboration (2017). Observation of a large-scale anisotropy in the arrival directions of cosmic rays above 8×10^{18} eV. *Science*, 357, 1266–1270. doi:10.1126/science.aan4338 · arXiv:1709.07321

52. Prager, B. J., Ransom, S. M., Freire, P. C. C., Hessels, J. W. T., Stairs, I. H., et al. (2017). Using long-term millisecond pulsar timing to obtain physical characteristics of the bulge globular cluster Terzan 5. *ApJ*, 845, 148. doi:10.3847/1538-4357/aa7ed7 · arXiv:1612.04395

53. Sanderson, R. E., et al. (WFIRST Astrometry Working Group) (2019). Astrometry with the Wide-Field Infrared Space Telescope. *JATIS*, 5, 044005. doi:10.1117/1.JATIS.5.4.044005 · arXiv:1712.05420

54. Socas-Navarro, H., et al. (2021). Concepts for future missions to search for technosignatures. *Acta Astronautica*, 182, 446–453. doi:10.1016/j.actaastro.2021.02.029 · arXiv:2103.01536

55. Soltis, J., Casertano, S., & Riess, A. G. (2021). The parallax of ω Centauri measured from Gaia EDR3. *ApJL*, 908, L5. doi:10.3847/2041-8213/abdbad · arXiv:2012.09196

56. Stappers, B., & Kramer, M. (2016). An update on TRAPUM. *PoS*, MeerKAT2016, 009. doi:10.22323/1.277.0009

57. Strader, J., Chomiuk, L., Maccarone, T. J., Miller-Jones, J. C. A., Seth, A. C., Heinke, C. O., & Sivakoff, G. R. (2012). No evidence for intermediate-mass black holes in globular clusters: strong constraints from the JVLA. *ApJL*, 750, L27. doi:10.1088/2041-8205/750/2/L27 · arXiv:1203.6352

58. Swanson, T. (2026). The Macro Transcension Hypothesis: spinning black holes in dense stellar clusters as thermodynamic attractors for advanced civilizations, with Omega Centauri as an observational test bed. Preprint; omegacentauri.me/macro-transcension-hypothesis.html (Paper A)

59. Tarter, J. (2001). The search for extraterrestrial intelligence (SETI). *ARA&A*, 39, 511–548. doi:10.1146/annurev.astro.39.1.511

60. Tremou, E., Strader, J., Chomiuk, L., Shishkovsky, L., Maccarone, T. J., Miller-Jones, J. C. A., et al. (2018). The MAVERIC survey: still no evidence for accreting intermediate-mass black holes in globular clusters. *ApJ*, 862, 16. doi:10.3847/1538-4357/aac9b9 · arXiv:1806.00259

61. van der Marel, R. P., & Anderson, J. (2010). New limits on an intermediate-mass black hole in Omega Centauri. II. *ApJ*, 710, 1063–1088. doi:10.1088/0004-637X/710/2/1063

62. Wright, J. T., Mullan, B., Sigurdsson, S., & Povich, M. S. (2014). The Ĝ infrared search for extraterrestrial civilizations with large energy supplies. I. *ApJ*, 792, 26. doi:10.1088/0004-637X/792/1/26 · arXiv:1408.1133

63. Wright, J. T., et al. (2022). The case for technosignatures: why they may be abundant, long-lived, highly detectable, and unambiguous. *ApJL*, 927, L30. doi:10.3847/2041-8213/ac5824 · arXiv:2203.10899

64. Zocchi, A., Gieles, M., & Hénault-Brunet, V. (2019). The effect of stellar-mass black holes on the central kinematics of ω Cen: a cautionary tale for IMBH interpretations. *MNRAS*, 482, 4713–4725. doi:10.1093/mnras/sty1508 · arXiv:1806.02157

Companion paper: The Macro Transcension Hypothesis (Paper A — the speculative hypothesis this campaign tests, alongside conventional astrophysics).

AI assistance disclosure: drafting, citation verification, derivation checking, and figure preparation were performed with substantial assistance from a large language model (Claude, Anthropic), under the author's direction; the author reviewed and takes full responsibility for all claims, derivations, and references.