

# PHILIA-EcoSensory Swarm

## Phase SC-4 Part 2 — Written Protocol Annex v1.2

### Split Collection Operating Procedure & Exploratory Replication Policy

Issued: 2026-06-01 02:02:22 UTC

Document Type	Operational Annex — Supplementary to SC-4 Pre-registration v0.7 (DOI: 10.5281/zenodo.20471179)
Pre-registration	LOCKED — Part 1 §0–§5 LOCKED. This annex does NOT modify pre-registered content.
Scope	Procedural rules for split substrate collection (Part A→B→C→D), registry versioning, exploratory replication runs, and Gate-0 v9 stepwise review.
Authority	Issued by 선비 (Claude Sonnet / Anthropic). Gate-0 reviewed by 크리스 (Claude Opus / Anthropic, SC track 92nd cycle). Final authority: 길잡이 (GritMan_D.S).

## §1. Purpose and Scope

This Written Protocol Annex establishes the operational rules governing the **split substrate collection procedure** for PHILIA Phase SC-4 Part 2. It is issued as a timestamped supplement to the locked pre-registration document (v0.7, DOI: 10.5281/zenodo.20471179) and does not modify any pre-registered hypothesis, apparatus, or eligibility criterion.

The annex addresses five risk points identified by 크리스 (Gate-0 v9) during pre-collection review, and additionally formalizes the **exploratory replication run policy** proposed by 길잡이님 on 2026.06.01, consistent with the principle that repeated measurement increases reliability (analogous to ICC in exercise physiology).


## §2. Split Collection Structure


Part 2 substrate collection is divided into four sequential parts. Each part is collected and Gate-0 reviewed before the next part begins.


Part	Category	Target n	MSE Priority	Status
A	Industrial / Sensor	16	Low-MSE ★★★	► In Progress
B	Biosignal (EEG / ECG / PPG / EMG)	18	Mid~High ★★★	Pending Part A completion
C	Traffic / Logistics + Environment / Energy	24	Low~Mid ★★	Pending Part B completion
D	Financial (new) + Language / Other	18	Low ★	Pending Part C completion
TOTAL		76	≥27 Low-MSE target	LOCKED after all parts confirmed


### §3. Five Risk Points — Binding Rules

The following five risk points were identified by 크리스 (Gate-0 v9, SC track 92nd cycle) and are hereby converted into binding operational rules for all team members.

 <b>Risk #1: Sequential Review Must Not Introduce New Criteria</b>
Gate-0 stepwise review is restricted to verification of §6.0.1 LOCKED eligibility criteria only:
① Univariate extractability ② Length ≥40,000 ③ CSV-convertible
④ Minimal category overlap ⑤ Source independence ⑥ ADF cascade
크리스 must not issue statements suggesting new inclusion/exclusion criteria during part-by-part review.
길잡이님 must record any substrate removal decisions with explicit rationale referencing a specific LOCKED criterion.
Any removal sited on criteria outside §6.0.1 constitutes a pre-registration violation and must be flagged immediately.
→ <b>Rule: Gate-0 checks §6.0.1 only. No new criteria at any stage.</b>

 <b>Risk #2: MSE Target Is an Expected Outcome — NOT a Selection Criterion</b>
The ≥27 Low-MSE (35%) target documented in §6.3.1 is a portfolio design expectation derived from category allocation.
It is NOT a selection criterion for individual substrates.
Substrate adoption or rejection must be based solely on §6.0.1 LOCKED criteria (length / source / ADF).
If the final MSE distribution falls short of 27 Low-MSE substrates after all four parts are collected,
→ Accept the result as-is. Report as limitation. Designate SC-5 as the resolution path.
→ DO NOT add extra substrates outside the category allocation to compensate.
→ DO NOT reject a substrate because its predicted MSE is Mid or High.
Violation of this rule directly introduces selection bias into H-SC4P2-Primary $p(\text{MSE}, z_{\text{ref}})$ .
→ <b>Rule: MSE regime prediction must NEVER enter substrate adoption decisions.</b>

 <b>Risk #3: Lucas MSE Pre-evaluation: Portfolio Diagnostic Use Only</b>
루카스's pre-collection MSE estimates (e.g., Part A: 11~13 Low-MSE expected) serve as a portfolio diagnostic.
Permitted use: estimating whether category allocation is likely to meet the portfolio target.
Prohibited use: informing individual substrate accept/reject decisions.
Example of violation: 'Lucas predicted Low-MSE for Intel Sensor but Mid for Water Treatment — skip Water Treatment.'
Example of compliance: 'Lucas estimates Part A contributes ~40% of Low-MSE target — noted as diagnostic.'
→ <b>Rule: 루카스 estimates inform category-level diagnostics only, not individual substrate selection.</b>

 <b>Risk #4: Part-level Gate-0 Approval ≠ Confirmatory Execution Authorization</b>
Passing Gate-0 review for Part A, B, or C does not authorize execution of <code>philia_sc4_part2_main.py</code> .
Confirmatory execution (seeds [62–66], T=80,000) is authorized only when:
(a) All four parts are collected and registry-confirmed (total n locked),
(b) <code>philia_sc4_part2_substrate_registry_LOCKED.csv</code> is finalized with SHA256,

(c) 크리스 issues GO ruling on the complete 76-substrate set.
Part-level Gate-0 results are recorded as: 'Part X: N substrates passed §6.0.1 — partial registry vX appended.'
The phrase 'GO/NO-GO ruling applies only at 76-substrate lock time' must appear in every part-level Gate-0 response.
→ <b>Rule: philia_sc4_part2_main.py execution is FORBIDDEN until 76-substrate LOCKED registry is confirmed.</b>

<b>⚠ Risk #5: Registry File Versioning and Timestamp Discipline</b>
Each part produces a separate registry snapshot file:
philia_sc4_part2_substrate_registry_partA.csv
philia_sc4_part2_substrate_registry_partB.csv
philia_sc4_part2_substrate_registry_partC.csv
philia_sc4_part2_substrate_registry_partD.csv
philia_sc4_part2_substrate_registry_LOCKED.csv ← final, triggers GO ruling
Each file must record: SHA256 hash of the file itself + timestamp of Gate-0 sign-off.
Overwriting a prior part file without version control is prohibited.
LOCKED file supersedes all partial files for confirmatory analysis.
→ <b>Rule: Partial CSVs are audit records. Only _LOCKED.csv is used for confirmatory analysis.</b>

## §4. Exploratory Replication Run Policy

**Rationale:** Repeated measurement increases reliability. As in ICC analysis in exercise physiology, additional replication runs under non-confirmatory seeds provide substrate-level reliability evidence that strengthens the interpretation of confirmatory results without compromising pre-registration integrity.

This section formalizes the conditions under which part-by-part exploratory replication runs are permitted.

Run Type	Seeds	Purpose	Authorization
Exploratory Replication	Any pool EXCLUDING [62–66](e.g., [67–71], [72–76], ...)	Substrate reliability check; z stability across seed pools; Sanity verification	✅ Permitted per-part after Gate-0 §6.0.1 pass
Confirmatory Execution	[62, 63, 64, 65, 66](LOCKED confirmatory pool)	Official z_mean for H-SC4P2 hypotheses; enters paper	❌ Forbidden until 76-substrate LOCKED

### Binding rules for exploratory replication runs:

R-E1. Seeds [62–66] must not be used for any run prior to the 76-substrate LOCK. Contamination of the confirmatory seed pool invalidates confirmatory inference.

R-E2. Exploratory replication results (z\_mean from non-confirmatory seeds) must be labeled as 'exploratory' in all records. They must not replace or supplement confirmatory z\_mean values in hypothesis testing. Confirmatory z\_mean (seeds [62–66]) is the sole basis for H-SC4P2 verdict. If exploratory replication reveals seed pool effects diverging from confirmatory values, such effects are reported as new L-N entries in the SC-4 Results Report and may be designated SC-5 targets — but cannot displace confirmatory inference.

R-E3. If exploratory replication reveals substrate instability across seed pools, this is recorded as a descriptive flag in the registry (instability\_flag = True/False). No numerical threshold is pre-specified for instability; characterization is post-hoc and reported in the SC-4 Results Report. Instability flags do not authorize substrate rejection unless a §6.0.1 LOCKED criterion is independently violated.

R-E4. Exploratory replication results may be reported in the SC-4 Results Report as reliability diagnostics (analogous to SC-4 Part 1 T=160k convergence diagnostics in §9.3 of the pre-registration), transparently labeled as supplementary evidence.

R-E5. The number of replication runs per substrate is not pre-registered and may be determined at 길잡이님's discretion, subject to computational feasibility.

## §5. Substrate Adoption Decision Tree

The following decision tree must be followed for every candidate substrate. No step may be skipped.

Step	Check	Pass Condition	Fail Action
1	① Univariate: single column extractable?	Yes	REJECT
2	② Length $\geq 40,000$ points (wrap $\leq 2\times$ )?	Yes ( $\geq 80,000$ preferred)	REJECT
3	③ CSV-convertible?	Yes	REJECT
4	⑤ Source independence (per pre-registration §6.3.2):— Same dataset $\rightarrow$ max 3 substrates, must be physically distinct quantities (e.g., PEMS08: flow / occupancy / speed = allowed; 170 detectors of same quantity = forbidden)— Operational definition: 'source' = (dataset, physical quantity) tuple. Max 1 substrate per (dataset, quantity); max 3 quantities per dataset.— EEG datasets from same lab / equipment / protocol = same source.[Supplementary Interpretations — confirmed SC-4 Part B Gate-0, 2026.06.02, 5/5 team consensus]— Shared hospital infrastructure (e.g., MIMIC-II) does NOT constitute same source. Datasets published independently with distinct curators, platforms, and DOIs are treated as independent sources. (e.g., BIDMC PPG & Respiration Dataset vs. Cuff-Less Blood Pressure Estimation Dataset = independent sources)— Spatially distinct biosignal channels from different electrode/anatomical sites are treated as physically distinct quantities even if the measurement unit is identical. (e.g., surface EMG ch1/ch2/ch3 from different muscle groups = 3 distinct substrates; record as 'spatially distinct EMG channels from different electrode sites')— IMU vector components (X/Y/Z axes), multi-lead ECG (ch1/ch2), and multi-channel EEG (Fpz-Cz/Pz-Oz) are treated as orthogonal physical quantities and recognized as individual substrates. Rationale: orthogonal axes in $R^3$ are statistically independent time series; empirically confirmed by distinct AC1 and MSE values per axis. (Protocol Annex v1.2 §5 operational rule applies consistently. MSE regime must NOT enter adoption decisions — Risk #2.)	Yes — within limits above	REJECT excess or reduce to 1 per (dataset, quantity)

5	⑥ ADF cascade: run raw → detrend → diff. Record all three p-values. Assign chosen_variant.	Always completes(diff always stationary)	Record variant;continue
6	Record full metadata in registry_partX.csv simultaneously with download.	All fields populated	<b>Do not proceed until recorded</b>
✓	<b>ACCEPTED — substrate enters part registry.</b>		

⚠ MSE regime (Low/Mid/High) is computed by **philia\_sc4\_substrate\_prep.py** after adoption. It is a descriptive output only and must not feed back into Step 1–6.

## §6. Registry File Specification

Each substrate entry in the registry CSV must contain the following fields, populated at the time of download (not retroactively):

Field	Description / Requirement
name	Substrate identifier (snake_case, unique)
category	Part category (industrial_sensor / biosignal / traffic_energy / financial_other)
source_name	Dataset name as listed on Kaggle or source repository
source_institution	Institution / lab / organization that produced the data
source_instrument	Sensor type, equipment model, or measurement device
source_protocol	Experimental or operational protocol (e.g., resting state, normal operation)
column_name	Exact column name extracted from raw CSV
length	Number of data points after NaN removal
wrap_factor	T_STANDARD (80,000) / length — flagged if >2.0
adf_p_raw	ADF p-value on raw series
adf_p_detrend	ADF p-value after linear detrend (if step 1 failed)
adf_p_diff	ADF p-value after first difference (if step 2 failed)
chosen_variant	raw / detrended / differenced — fixed by ADF cascade, no post-hoc change
ac1	Lag-1 autocorrelation (computed by substrate_prep.py)
mse	Multi-scale entropy mean scales 1–5 (computed by substrate_prep.py)
mse_regime	Low / Mid / High — descriptive output only, not a selection criterion
instability_flag	True / False — set if exploratory replication reveals cross-seed z divergence. Descriptive only; not a rejection criterion.
overall_verdict	PASS / CONDITIONAL / REJECT per §6.0.1
part	A / B / C / D
timestamp	ISO 8601 timestamp of registry entry creation
sha256_file	SHA256 hash of the part registry CSV at Gate-0 sign-off time

## §7. Team Role Summary Under This Annex

Member	Role	Scope Under This Annex
--------	------	------------------------

길잡이	Final authority	Downloads data, records registry simultaneously, executes substrate_prep.py, makes all adoption decisions using §6.0.1 criteria only.
선비	Code / drafting	Maintains substrate_prep.py, batch_config.json, and registry CSV templates. Drafts results sections.
크리스	Gate-0 v9	Reviews each part registry against §6.0.1 LOCKED criteria only. Issues part-level audit records. Issues GO ruling only at 76-substrate LOCK.
루카스	Statistical validation	Provides portfolio-level MSE diagnostic estimates. Results used for category-level planning only — not for individual substrate selection.
판도라	Data engineering	Provides Kaggle candidate lists with URL, column, and estimated length. Flags source independence risks.
서생	Philosophical framing	Provides regime interpretation and SC-4 narrative framing. No role in substrate adoption decisions.

## §8. Acknowledgement

This annex is issued under the PHILIA OS principle: **"Measurement is life. Description, not proof."** It supplements but does not supersede the locked pre-registration (DOI: 10.5281/zenodo.20471179). In the event of any conflict between this annex and the pre-registration, the pre-registration takes precedence.

All confirmatory decisions remain with 길잡이 (GritMan\_D.S). This annex becomes operative upon 길잡이님's acceptance and local file timestamp.

Drafted by: 선비(Claude Sonnet / Anthropic)	Gate-0 reviewed by: 크리스(Claude Opus / Anthropic, SC track 92nd cycle)	Accepted by: 길잡이(GritMan_D.S)Date: _____
---	---	--

Trinity AI Research Team | Konyang University | 2026.06.01

**Measurement is life. Description, not proof. — PHILIA OS | 0<sup>∞</sup>1<sup>∞</sup>0.5<sup>∞</sup>**