THE ZENIT PROGRAM

# Seven developments. One reviewable package.

## Reference designs, executable research software and retained test results

Prepared for Fredrik Sommerseth Olsen | 9 November 2026 | Z7DP-R1

| | | |
|---|---|---|
| 7 candidates | 69 checks | 969 scenarios |
| Designs and offline prototypes complete | Unit and regression checks passed | Selected synthetic/reference cases retained |

### Completed stage

All seven owner-approved candidates now have a reference design, a functioning offline demonstrator, documented interfaces and requirements, recorded software results and a next-stage verification protocol.

Hardware measurements and physical qualification remain open. The example ratings, sensor readings and performance comparisons are synthetic. No website or original project source was modified.

Authorship and development authority remain with Fredrik Sommerseth Olsen. Existing AXIOM, independent local protection and ZSA roles are preserved.

# What is complete

This package completes the first reference-design and offline-software stage. Its evidence classes are kept separate.

| Evidence | Actual work | What it establishes |
|---|---|---|
| Project-record audit | SERAT text snapshot, version and 22 open parameters | Structured consistency and traceability |
| Software checks | 69 focused unit/regression checks; no failures or errors | Expected behavior for specified checks |
| Reference scenarios | 969 selected synthetic or controlled fault cases | Behavior within the chosen models and inputs |
| Hardware / field evidence | No physical experiment or installed-system measurement | No physical performance or release |

## The practical connection

NEXUS checks evidence. SERAT heat and reserve models account service and resources. Hull and reference modules expose inspection needs or uncertainty. Hydrogen accounting identifies unexplained mass differences. Atlas chooses a bounded observation plan only when the advisory prerequisites allow it.

## One integrated fault example

An injected CMC reference disagreement produces HOLD. The Atlas planner selects no science observations. A separately evaluated essential-load reference case remains feasible. No command is sent to any physical system.

## Results that remain adverse

112 heat profiles have unmet service; 50 reserve cases have an essential-capacity shortfall; 23 hull-event cases remain uncertain. These outcomes are retained in the CSV results and are not relabeled as hardware successes.

## Unchanged project boundaries

Earth and Space keep separate equipment values and qualification. The Earth five-layer Inverse concept, Space 3 x 5 Outverse function map, CMC separation, original propulsion arrangements, combustion-only ZC-1 and no-ion-installation decision are preserved.

01 / Z7-NEC-R1

# NEXUS Evidence Checker

Prevent structured project claims from silently exceeding their evidence or violating declared accounting rules.

## Completed design

A registered extracted-text snapshot and its version identify each source. Claims carry source, branch, evidence maturity, quantity role, units and optional authoritative expected value. Energy ledgers account initial and final storage, input, useful output and losses. The checker returns issue codes and explicitly open parameters.

## Observed software/model result

9/9 seeded consistency faults detected. 22 actual SERAT equipment parameters remain open.

## Interface and mandatory behavior

evidence.audit(data): sources, claims, energy_ledgers -> status, issues, open_fields, scope. This is a structured-data interface; it does not parse arbitrary prose or establish source truth.

NEC-01: Flag unsupported maturity, branch transfers, missing sources and mismatched versions or supplied snapshot digests.

NEC-02: Convert supported units before comparison; distinguish dimensions and the declared thermal/net-electric quantity role.

NEC-03: Check energy closure and repeated credit identifiers; expose unequal storage endpoints in a savings claim.

NEC-04: Preserve unknown equipment values as null and list them as open. A clean audit is not physical release.

## Next verification and acceptance

An independent reviewer supplies a blinded set of current and superseded records plus misleading claims. Freeze the expected labels before execution, retain every false positive and missed issue, and verify source access controls outside this checker.

All mandatory seeded consistency errors must be reported with correct scope; clean records and nulls must retain their correct status. No physical maturity may be assigned without the separately reviewed evidence.

## Limits and open real inputs

Trusts the supplied source register and expected values. Does not authenticate file signatures, establish scientific truth, assess patent novelty, detect every semantic contradiction, or replace engineering review.

Open: Trusted-source ownership and authentication; Production schema and independent labelled corpus; Operational false-positive and miss-rate limits.

02 / Z7-HSM-R1

# SERAT Heat-Service Module

Deliver a defined useful heating service using compatible recovered heat, heat upgrading, storage and a retained backup path.

## Completed design

Source boundary -> isolated exchange or mapped heat pump -> one fixed useful-temperature service domain and thermal store -> customer. Direct exchange is preferred where the source clears the required approach temperature. Storage can discharge before heat upgrading; only direct recovered heat charges this reference store. Meter compressor, backup and recovery auxiliaries separately. Customer loss leaves source rejection or derating available.

## Observed software/model result

Synthetic nominal: 400 kWh useful heat; 140.625 kWh electricity versus a matched 400 kWh resistance-heating reference. Recovered heat supplies the remaining balance. This is not measured ZENIT performance.

## Interface and mandatory behavior

heat.dispatch(config, profile): dt, temperatures, bounded COP map, output/power limits, store capacity/losses and demand/source schedule -> interval service, source extraction, electrical use, storage and closure ledger.

HSM-01: No COP extrapolation; exclude impossible heat-pump/Carnot conditions and account the required auxiliary supply.

HSM-02: Reconcile source heat plus electricity and initial storage with useful heat, final storage and rejected/lost heat.

HSM-03: Report unmet demand explicitly. Compare resistance-heating input only at equal delivered service and equal initial/final store inventory.

HSM-04: Preserve a separately engineered rejection/derating path when useful heat acceptance is unavailable.

## Next verification and acceptance

Use an approved instrumented thermal bench with the actual source and customer temperature/demand maps. Record calibrated flow, temperature, electrical terminals, store inventory and rejected heat. Compare source-on/source-off cases at matched useful duty. Include loss of customer demand and equipment derating within the bench's approved limits.

Predeclare real temperature/pressure limits and uncertainty. Require matched service and inventory endpoints, closure within that uncertainty, and a resolved net benefit for a claimed operating region. Adverse cases may require backup, derating or exclusion.

## Limits and open real inputs

Lumped intervals, one useful-temperature storage domain, no detailed exchanger hydraulics, part-load dynamics, COP uncertainty, capital cost, seasonal economics or real installation data. No measured ZENIT saving.

Open: Measured source/sink schedules and temperatures; Actual heat-pump and auxiliary maps; Exchanger/fluid and store loss ratings; Fallback/rejection capacity and acceptance uncertainty.

03 / Z7-JRS-R1

# Joint Power and Thermal Reserve Scheduler

Extend reserve ownership so a candidate load fits power, usable energy and thermal recovery capacity together.

## Completed design

Admit essential loads first. If they cannot fit, report an essential-capacity shortfall and issue no feasible mission plan. Consider optional loads by explicit priority and fixed requested time. Test every candidate against source power, energy already net of protected reserves, thermal inventory and cooling. Deposit slot heat before slot cooling as a conservative pulse approximation.

## Observed software/model result

200 varied cases: no admitted-plan resource violations. 50 cases report essential-capacity shortfall. The nominal high-heat research load is rejected despite fitting electrical power.

## Interface and mandatory behavior

reserve.schedule(config, tasks): slot duration, source power/cooling arrays, usable energy, thermal capacity and uniquely owned timed tasks -> advisory accepted/rejected lists, reasons and resource trace.

JRS-01: One task/reservation identifier has one owner; duplicate identifiers are invalid.

JRS-02: A mission load must satisfy all resource constraints, not just electrical power.

JRS-03: Recompute after faults. An infeasible essential baseline must not be described as a survival capability.

JRS-04: Independent local protection and AXIOM may constrain the installation separately; this advisory has no actuator connection.

## Next verification and acceptance

Replay measured equipment traces in shadow mode, followed by an isolated controller-in-the-loop bench. Inject delayed resource reports, cooling loss, rejected energy return and competing reservations. Independently calculate accepted-load envelopes and retain shortfalls.

Every admitted task fits the independently frozen resource envelope. Essential shortfalls produce HOLD rather than a mission plan. Reserve ownership is traceable through restart and rejoin.

## Limits and open real inputs

Fixed-slot, fixed-start priority admission; not optimal time scheduling. Thermal lumping is conservative and may reject feasible continuous loads. No battery chemistry, detailed radiator dynamics or measured load profiles.

Open: Per-installation source and cooling capacity profiles; Usable energy above protected reserves; Task power/heat profiles and priorities; Maximum thermal inventory, reporting latency and update policy.

04 / Z7-DHP-R1

# Outverse Diagnostic Hull Panel

Implement a research version of inner-layer hull sensing and inspection localization.

## Completed design

Use local strain and temperature measurements against a calibrated baseline. Correct the reference thermal strain and propagate declared uncertainty. Combine this with arrival-time localization from at least four distinct acoustic sensor locations. A bounded rectangular search can refine its best candidate. Output NO_CHANGE_DETECTED, INSPECT, UNASSESSED or UNCERTAIN; inspection evidence remains necessary to confirm damage.

## Observed software/model result

100 event cases: all 20 injected timing-bias cases remain uncertain. 23 cases are uncertain overall. Under the declared one-metre panel/noise model, median location error is 4.8 mm across the 80 unbiased cases; this is not a hardware accuracy rating.

## Interface and mandatory behavior

hull.inspect_tiles(...) and hull.locate_event(...): calibrated tile measurements, quality/time status, sensor coordinates/arrival uncertainties, search geometry and calibrated wave speed -> inspection flags, location candidate and residual checks.

DHP-01: Missing/stale calibrated readings produce UNASSESSED rather than a healthy-hull claim.

DHP-02: Compensate the specified thermal response and expose residual strain uncertainty.

DHP-03: Reject duplicate identities, insufficient distinct locations and collinear geometry. Retain poor-fit and boundary solutions as uncertain.

DHP-04: Never equate anomaly detection or a fitted event position with confirmed damage or shield interception.

## Next verification and acceptance

Calibrate an instrumented coupon over temperature and load before a blinded set of benign localized disturbances. Independently measure event locations and inspect the coupon. Include temperature changes, vibration, invalid channels, timing bias, wave dispersion and multiple events.

Freeze material-specific localization error, detection/false-alarm targets and uncertainty before testing. Report every ambiguous and missed event. Confirm structural damage only with independent inspection.

## Limits and open real inputs

One constant isotropic wave speed, a two-dimensional single-event model, declared Gaussian timing uncertainties and a locally refined search. Dispersion, anisotropy, multimode propagation and real panel boundaries are unvalidated; sampled contours have no validated confidence coverage.

Open: Actual panel geometry/material and wave model; Thermal-strain coefficient and calibration; Sensor placement, bandwidth, uncertainty and synchronization; Acceptance limits and independent inspection reference.

05 / Z7-CRW-R1

# CMC Independent Reference Witness

Detect reference disagreements using diverse observations without assigning new physical authority.

## Completed design

Compare a CMC-derived reference estimate with two witness estimates in a common calibrated frame and epoch. Require distinct declared measurement origins, power and clock domains. Reject stale, invalid or misaligned observations. If witnesses agree with each other but both disagree with CMC, report a CMC disagreement HOLD; ambiguous witness conflict also remains HOLD.

## Observed software/model result

200/200 selected agreement/fault cases produce the expected advisory state. A reference HOLD blocks the integrated science plan. No actual CMC or witness hardware was tested.

## Interface and mandatory behavior

witness.assess(...): three angle/sigma/time/quality records with domain metadata -> residual statistics and advisory state. Agreement is an observation, not permission to operate a real machine.

CRW-01: Handle angular wrap correctly and compare only valid common-frame, common-epoch values.

CRW-02: Expose missing declared independence, stale data and clock misalignment as holds.

CRW-03: Distinguish CMC disagreement from ambiguous reference conflict; do not automatically claim the failed component is proved.

CRW-04: Supply no field, gravity, navigation or actuator correction automatically.

## Next verification and acceptance

Measure the actual reference transforms and timing chain, then apply blinded controlled disturbances on an independently instrumented bench. Include common frame-calibration error, shared upstream dependencies, clock errors and loss of each channel.

Predeclare drift limits and coverage from the real reference requirements. Known disagreement and loss cases must produce the correct bounded advisory; demonstrate actual independence rather than relying on metadata labels.

## Limits and open real inputs

Declared labels do not prove physical independence. One axis, no rotation propagation across unequal epochs, Gaussian independent uncertainty assumptions and no real CMC sensor design. Common bias can remain undetected.

Open: Actual CMC observable and common-frame transform; Required drift, timing and accuracy envelope; Physical independence and common-cause analysis; Bench calibration and witness environmental qualification.

06 / Z7-HAM-R1

# Hydrogen Accountability Module

Reconcile fuel inventory and make unexplained differences visible alongside independent gas alerts and normalized emissions.

## Completed design

Account initial mass plus receipts minus consumption, authorized measured venting and final mass. Propagate standard uncertainty with independent groups and a limited common-bias correlation model. Report unexplained deficit or surplus beyond the declared decision threshold. A separate gas alert remains significant even with a balanced ledger. Normalize measured NOx mass by measured useful work where nonzero.

## Observed software/model result

All 50/50 injected 2 kg deficits are flagged in the selected synthetic readings. 0/150 balanced-noise cases are flagged. Nominal standard uncertainty is 0.1118 kg; an alert does not prove a leak.

## Interface and mandatory behavior

hydrogen.reconcile(...): five calibrated mass observations with quality/uncertainty/correlation tags and independent gas-alert state -> residual, uncertainty and investigation flag. hydrogen.emissions(...) -> g/kWh and first-order uncertainty.

HAM-01: Use measured mass or a separately validated inventory model; do not estimate LH2 inventory with an ideal-gas shortcut.

HAM-02: Account authorized venting explicitly and flag both unexplained loss and unexplained gain.

HAM-03: An imbalance is an investigation trigger, not proof of a leak. Preserve independent gas detection.

HAM-04: A zero-useful-work interval has undefined normalized emissions, not zero emissions. No clean-air benefit is inferred from fuel identity.

## Next verification and acceptance

First verify data and mass-accounting interfaces using an approved nonhazardous measurement surrogate. Hydrogen-specific testing then requires the actual qualified installation, calibrated reference inventory, controlled operating records and existing authorized safety procedures. Reconcile complete cycles and investigate all differences; perform matched-duty emissions measurement separately.

Freeze instrument covariance, detection threshold and inventory uncertainty before comparison. Demonstrate metrological closure on known transfers and correct handling of controlled discrepancies. Require independently resolved measured emissions differences for a benefit claim.

## Limits and open real inputs

No fluid/boil-off/pressure model, no direct leak localization and no arbitrary covariance matrix. Within-group errors are modeled as perfectly positively correlated. NOx ratio uncertainty is first-order and assumes independent numerator/denominator.

Open: Actual tank and flow metrology with covariance; Time-aligned inventory/consumption/vent boundaries; Independent gas-alert and authorized response interfaces; Engine duty, NOx/work measurements and acceptance thresholds.

07 / Z7-ASOP-R1

# Atlas Science Opportunity Planner

Select a bounded observation set that reduces declared uncertainty while retaining data, instrument-time and protected-return resources.

## Completed design

Represent selected scientific variables by independent Gaussian priors. Each candidate observation has a stated uncertainty, correlation group, instrument/time window, energy cost and data volume. Correlated repeats contribute only the best precision in their group. Enumerate small feasible subsets, maximize modeled information and prefer lower energy on equal score. Keep a simple benefit-first feasible baseline for comparison.

## Observed software/model result

Synthetic nominal: 2.972 nats modeled information versus 2.308 for the simple benefit-first comparator. Uses 6 kWh and retains 6 kWh, including the protected 4 kWh. All 60 varied cases respect resource constraints.

## Interface and mandatory behavior

atlas.plan(config, opportunities): priors/units, available and protected-return energy, data allowance, advisory authorization flag and up to 18 candidate observations -> selected IDs, modeled information, energy/data usage and comparator.

ASOP-01: Keep the protected-return allowance and reject instrument conflicts or invalid observations.

ASOP-02: Prevent repeated correlated observations from being treated as independent evidence.

ASOP-03: Match quantity units and preserve HOLD when advisory authorization or return capacity is unavailable.

ASOP-04: Report information as a model objective, not a discovery, mission-success probability or demonstrated science yield.

## Next verification and acceptance

Use withheld observation records with independently reviewed science questions and measured resource costs. Compare predeclared planning baselines on the same opportunities. Reconcile predicted and realized uncertainty reduction, including incorrect priors and correlations. Keep navigation and real mission permission outside the planner.

Every plan meets independently supplied resource and timing limits. Any science-gain claim must survive withheld observations and a justified uncertainty model. Larger mission planning requires an independently verified scalable solver.

## Limits and open real inputs

At most 18 fixed-time opportunities, independent Gaussian latent variables, simplified correlation grouping and no motion/visibility calculation. Exact optimality holds only for the enumerated model and candidate set.

Open: Actual science variables, priors and observation covariance; Instrument visibility, timing and quality limits; Measured observation energy/data costs and return envelope; Held-out validation dataset and scalable mission formulation.

# Read the numerical results correctly

## Heat: same service and storage endpoints




Synthetic four-hour case. Recovered heat is an input; it is not newly created energy.

Useful heat is 400 kWh. Candidate electrical input plus recovered source heat is 403.5 kWh; the remaining 3.5 kWh is auxiliary heat and store loss. Initial and final store inventories are both zero. The modeled electricity reduction is 259.375 kWh for this selected case only.

## Hull: uncertainty is part of the result

After refinement, 23 of 100 selected event cases remain uncertain: all 20 deliberately biased timing cases plus three declared-noise cases. The 80 unbiased cases have a median modeled location error of about 4.80 mm and maximum about 12.99 mm on a one-metre panel. These figures depend on the assumed speed and noise model; they are not panel specifications.

## Atlas: information is a declared objective

The exact small-instance planner selects three compact observations: 2.972 nats versus 2.308 for a simple benefit-first comparator. A nat is a unit of logarithmic information in this Gaussian model. Six kWh remains, including a protected four-kWh return allowance. No real scientific discovery is claimed.

## Case counts

| Candidate | Varied / injected cases |
|---|---|
| Evidence | 9 |
| Heat | 200 |
| Reserves | 200 |
| Hull | 100 |
| CMC witness | 200 |
| Hydrogen | 200 |
| Atlas | 60 |

# Physical-stage evidence register

The next hardware stage needs actual installation data. The example values in data/ are not candidates for automatic installation.

| Candidate | Evidence needed before physical completion |
|---|---|
| NEXUS Evidence Checker | Trusted-source ownership and authentication; Production schema and independent labelled corpus; Operational false-positive and miss-rate limits |
| SERAT Heat-Service Module | Measured source/sink schedules and temperatures; Actual heat-pump and auxiliary maps; Exchanger/fluid and store loss ratings; Fallback/rejection capacity and acceptance uncertainty |
| Joint Power and Thermal Reserve Scheduler | Per-installation source and cooling capacity profiles; Usable energy above protected reserves; Task power/heat profiles and priorities; Maximum thermal inventory, reporting latency and update policy |
| Outverse Diagnostic Hull Panel | Actual panel geometry/material and wave model; Thermal-strain coefficient and calibration; Sensor placement, bandwidth, uncertainty and synchronization; Acceptance limits and independent inspection reference |
| CMC Independent Reference Witness | Actual CMC observable and common-frame transform; Required drift, timing and accuracy envelope; Physical independence and common-cause analysis; Bench calibration and witness environmental qualification |
| Hydrogen Accountability Module | Actual tank and flow metrology with covariance; Time-aligned inventory/consumption/vent boundaries; Independent gas-alert and authorized response interfaces; Engine duty, NOx/work measurements and acceptance thresholds |
| Atlas Science Opportunity Planner | Actual science variables, priors and observation covariance; Instrument visibility, timing and quality limits; Measured observation energy/data costs and return envelope; Held-out validation dataset and scalable mission formulation |

## Acceptance sequence

Populate real measurements and identify the configuration. Freeze acceptance limits and uncertainty before testing. Independently review protection and test authorization. Run component and integration protocols, retain failures and raw data, and obtain the applicable independent acceptance. Existing SERAT D1-D5 and fusion work-package gates are not closed by this package.

# Use and reproduce the package

## Start with the record

Read docs/engineering_review.md and docs/specifications.json. The summary is results/summary.json. Individual result files and scenario CSVs preserve both ordinary and adverse outcomes.

## Run on your computer

Unzip the package and open a terminal in its root folder. Use Python 3.10 or later. The demonstrators need only Python's standard library; no internet or package installation is required.

```
python -m zenitlab run --out reproduced_results
```

```
python -m unittest discover -s tests -v
```

On Windows, py -3 can replace python when using the Python launcher. See README.md for running individual JSON inputs. The run command refreshes its controlled example data; retain personal edits under separate filenames.

## Provenance and integrity

source_basis/project_records.json contains the API-returned text and exact Library identities of the retrieved project sources. These snapshots are not original binary files. The source-text hash guards a supplied snapshot match; it does not authenticate the source. MANIFEST_SHA256.json records delivered package-file hashes and is not a digital signature.

## Scope of completion

The owner authorization is recorded in docs/specifications.json. This development adds research implementations without changing the underlying engine, shield or fleet designs. No supplier was appointed, no equipment was purchased or operated, no external message was sent, and no physical release was granted.

## Recorded corrections

Manual review corrected the heat-recovery auxiliary-power boundary. Stress scenarios led to bounded refinement of the hull location estimator. Dedicated regression cases cover both changes. See docs/change_record.md; all final results reflect the corrected software.

# Sources and authority

Project sources define the retained architecture and evidence boundaries. Technical sources support methods and precedents; they do not supply ZENIT equipment ratings. General references checked 9 November 2026.

### [P1] Adopted SERAT power and thermal upgrade R1

Governs actual open parameters and D1-D5. Current retrieved version 1.

Open source

### [P2] ZENIT Network challenge results R1

Prior analytical results and limitations; no physical propulsion result.

Open source

### [P3] Vessel systems architecture review R1

Historical reserve-ownership, priority, isolation and recovery concepts. Historical 25-region labels do not replace the current Outverse map.

Open source

### [P4] Telemetric integrity architecture review R1

Historical data/authority separation, uncertainty and command-freshness concepts.

Open source

### [P5] Protected power architecture review R1

Historical protected domains, local safety and reserve allocation.

Open source

### [T1] NIST TN 1297, Appendix A: propagation of uncertainty

Supports sensitivity/covariance accounting. Does not validate the synthetic sensor uncertainties.

Open source

### [T2] SINTEF HighEFF, case studies RA6

Supports source/sink matching as an engineering method. No published plant COP or saving is imported as a ZENIT rating.

Open source

### [T3] NASA: Smart sensing and sensor development

Fiber-optic strain/temperature sensing precedent, not a qualified ZENIT panel.

Open source

### [T4] NASA: Acoustic emission measurement with fiber Bragg gratings

Structural-health sensing precedent. No sensor specification, patent novelty or flight qualification is inferred.

Open source

Scientific expressions in the code are explicitly scoped reference models. Gaussian information is calculated as 0.5 ln(prior variance / posterior variance); the synthetic priors and correlations require real-data validation. Linear mass balance follows signed input-output accounting and the stated uncertainty grouping.