POLITECNICO DI TORINO
Repository ISTITUZIONALE

Ecosystem-Aware Design: Formalizing Programmatic Fragility in Early-Phase ECLSS Sizing for Human Lunar Missions

Original
Ecosystem-Aware Design: Formalizing Programmatic Fragility in Early-Phase ECLSS Sizing for Human Lunar Missions / Rimani, J., Luccisano, G., Narducci, G., Fusaro, R., Viola, N.. - ELETTRONICO. - (2026). (77th International Astronautical Congress (IAC 2026) Antalya (Türkiye) 5-9 October 2026).

Availability:
This version is available at: 11583/3017539 since: 2026-11-04T20:50:51Z

Publisher:
International Astronautical Federation (IAF)

Published
DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

20 November 2026

Ecosystem-Aware Design: Formalizing Programmatic Fragility in Early-Phase ECLSS Sizing for Human Lunar Missions

Jasmine Rimani[a, *], Giacomo Luccisano[b, *], Giuseppe Narducci[b], Roberta Fusaro[b], Nicole Viola[a]

[a] *Department of Management and Production Engineering (DIGEP), Politecnico di Torino, 10129 Turin, Italy.* nicole.viola@polito.it (N.V.).
[b] *Department of Mechanical and Aerospace Engineering (DIMEAS), Politecnico di Torino, 10129 Turin, Italy.*

[*] Corresponding Authors: jasmine.rimani@polito.it (J.R.); giacomo.luccisano@polito.it (G.L.).

Abstract

Early-phase ECLSS design links resource performance to equipment mass and cost, but implementation also depends on suppliers, test facilities and integration capability. This paper examines those dependencies through the HUMANS terrestrial analog at Politecnico di Torino and a reference human lunar lander. The paper compares integrated and distributed air-processing architectures using a common membrane-water function; an open lander provides a physical reference. Resource sizing, development/testing cost breakdowns and TRIS maturation paths establish the engineering comparison. An ecosystem-readiness profile then records the required capabilities, supporting evidence, remaining work and structural dependencies for each architecture and milestone. The profile uses work cost and schedule margin as figures of merit and retains unresolved requirements explicitly. Conditional schedules include procurement, configuration adaptation, shared facilities and integrated acceptance. HUMANS procurement follows the end-2027 structure target, with readiness and testing completion required in Q1 2029. The scenarios identify how implementation constraints can change a design trade, while the available evidence remains insufficient to confirm delivery dates or configuration-specific procurement advantages. HUMANS is also examined as a potential setting for establishing process knowledge, integration practice and operating capability relevant to later flight development. A proposed evidence-transfer register distinguishes what ground testing could establish from the flight-specific verification still required. A POMDP extension connects the profile to sequential testing and development decisions, without aggregating readiness into a score; numerical policies require additional probability evidence. The contribution is a traceable connection from ECLSS architecture to its implementation pathway and evidence needs.

Abbreviations

ACLS	Advanced Closed Loop System
ECLSS	Environmental Control and Life Support System
ESA	European Space Agency
HLS	Human Landing System
HUMANS	Humans in the loop Moon and Mars Analog for Sustainability
MER	Mass Estimating Relationship
PoC	Proof of Concept
TRIS	Technology Roadmapping Strategy
TRL	Technology Readiness Level

1 Introduction

Environmental Control and Life Support Systems (ECLSS) are essential to human spaceflight, maintaining a habitable environment and providing resources such as air and water. Programme-level decisions shape this capability, such as which technologies to mature, which suppliers to develop, and which test facilities to fund, often before flight hardware is specified. Strengthening European ECLSS capability supports Europe's ambition toward a new and daring human exploration mission. Accordingly, ESA's exploration roadmap lists life support among the required enabling capabilities [1, 2].

In addition to assessing the technical feasibility of a new human-spaceflight mission or design concept, it is essential to evaluate the industrial capabilities available to deliver it. A concept may satisfy its technical requirements on paper, yet remain difficult to implement if the required components, materials, testing facilities, or specialist expertise are unavailable within the relevant industrial base. Reliance on suppliers outside the region [1] can introduce additional procurement, coordination, logistics, and export-control burdens. These dependencies may increase cost, lengthen lead times, and create schedule risks. If a critical capability does not exist and must be developed through a new company or a dedicated investment programme, the project also becomes more exposed to political and funding discontinuities. An industrial-capability assessment can therefore support decision-making by programme managers, agen-

[1] Region is used here to denote a geographically bounded industrial area within which organisations, infrastructure, skills, and supply-chain relationships can be mobilised to support the development, production, qualification, and maintenance of the system under consideration. Its boundaries are defined functionally rather than exclusively by administrative divisions and may therefore extend across neighbouring subnational areas or national borders.

cies and funding bodies in two complementary ways.

First, it can identify strategic gaps where targeted investment could help establish or strengthen the required capabilities. Second, it can inform technology selection by favouring solutions that can be supplied, integrated, qualified, and maintained within the programme's budget and schedule constraints. Examples of these possible disruptions include:

- A regenerative life-support concept may offer superior mass and consumables performance, but it may be a poor near-term choice if no regional supplier can produce and qualify its sorbent beds, catalytic reactors, membranes, pumps, sensors, or control software at the necessary reliability level. The system may then face long procurement lead times or dependence on non-European sources.
- A supplier may be able to manufacture a component, but programme delivery still depends on access to facilities for vibration, thermal-vacuum, EMC, contamination-control, leak, or life-cycle testing. Missing test infrastructure can become a critical schedule driver even when the design itself is mature.
- Pressure vessels, high-pressure oxygen hardware, non-metallic materials compatible with oxygen-rich environments, and certain additive-manufacturing processes require specialised design, manufacturing, and certification expertise. Limited access to qualified suppliers can constrain the design space or require additional qualification campaigns.
- If a programme depends on creating an entirely new industrial capability, its schedule may extend beyond a funding cycle or ministerial mandate. A later change in priorities can leave the capability incomplete, increasing the risk that the mission architecture is delayed, descoped, or cancelled.

This study originates from a specific programme need rather than from a general interest in regional industrial analysis. Politecnico di Torino is developing the HUMANS facility, a ground-based research infrastructure intended to support testing of human-spaceflight systems, operations, and enabling technologies. Its ECLSS activities include characterisation under realistic metabolic loads and relevant atmospheric conditions, over multi-day to multi-week test durations. The facility officially entered development in February 2026. The project target is readiness in Q1 2029, with equipment integration and testing completed. The structure is the identified programme blocking item: it is targeted for availability by the end of 2027, followed by equipment purchase, integration and testing. These dates and this sequence are author-provided project inputs, not literature-derived lead times or confirmed delivery records. Because the facility is being developed from the ground up, ECLSS architecture and component selection must account for this sequence as well as technical performance. In this context, supplier availability is a design variable. Before issuing procurement calls, the project needs reasonable evidence that a sufficiently broad set of potential suppliers, ideally across Europe, can provide the necessary hardware, services, expertise, and qualification support at an acceptable maturity level and within the expected budget. This reduces the risk of defining a technically attractive system whose implementation is delayed or made unaffordable by a limited or inaccessible supply base.

Therefore, our research intent can be summarised as: *How can industrial availability and supply-chain uncertainty be represented during the early design of an ECLSS or human-spaceflight system, to assess whether technically feasible solutions can be delivered within defined budget and schedule constraints, and which of the resulting gaps can be closed by programme action, such as supplier development, facility access or qualification campaigns, within the relevant funding cycle?*

The purpose is not to determine that one geographical region is universally "better" than another, nor to claim that industrial mapping produces a single preferred technical solution. Instead, it provides a structured way to examine practical deliverability under a cooperative European scenario: if the participating organisations intend to develop this system together, does the necessary industrial capability already exist, and where are the gaps or dependencies?

To address this question, the study combines an engineering assessment of candidate ECLSS architectures with an assessment of the industrial ecosystem required to implement them. For each architecture, the analysis identifies the relevant subsystem functions, component-level dependencies, technology maturity, supplier evidence, and the development, qualification, integration, and testing activities still required for deployment.

This information is used to identify capability gaps and to estimate the associated development effort, cost exposure, and schedule risk. Because evidence on supplier availability and capability maturity is often incomplete during early design, the methodology records these factors in an ecosystem-readiness profile. The profile connects supported, unresolved and insufficient capabilities to the work needed for implementation and to the associated cost and schedule figures of merit.

The method is not intended to predict whether a specific supplier will ultimately be selected or whether a particular development activity will succeed. Instead, it provides a structured basis for comparing candidate architectures according to both their technical performance and the readiness of the ecosystem required to realise them.

The methodology is applied to two case studies with different development conditions and objectives. The first case is HUMANS, as introduced before. It provides a ground-based environment for investigating habitat

operations and life-support technologies, including experimental work on air and water regeneration. The second case is a short-sortie lunar lander. This case represents a more demanding flight-development context, in which candidate architectures must satisfy tighter constraints on mass, reliability, safety, qualification, and mission readiness. It examines how the same assessment approach performs when technology maturation and industrial availability directly affect a flight system. Together, the two cases distinguish between a near-term terrestrial implementation need and a more ambitious exploration scenario. They allow the study to examine how differences in operating conditions, performance requirements, integration constraints, and acceptable development risk affect ECLSS architecture selection.

The paper makes three contributions.

- *Engineering.* A comparison of integrated and distributed air-processing architectures that follows each one from resource sizing to the implementation work it requires, in a ground facility and in a flight vehicle.
- *Method.* An ecosystem-readiness profile that keeps supported, unknown and documented-shortfall requirements apart, and a gap-to-action typology that links each unresolved condition to the class of programme action that would resolve it: information, investment or infrastructure.
- *Programme.* A reading of the results for programme and funding decisions that states the strength of the evidence behind each implication, and a proposed register through which a ground facility such as HUMANS can build capability for later flight programmes.

The results show how resource-recovery assumptions, technology maturity, supplier evidence, development work, and schedule constraints affect comparisons of candidate ECLSS architectures.

The outcome is not a fixed ranking of regions or technologies, but a structured basis for planning development under uncertainty and for identifying where targeted programme action would most reduce implementation risk.

Section 2 positions the study with respect to technology roadmapping and the literature on innovation ecosystems. Section 3 defines the two cases and Section 4 the assessment, including the readiness profile and the gap-to-action typology. Section 5 reports the engineering, cost and implementation results, Section 6 examines HUMANS as a public research infrastructure, Section 7 draws the implications for programme and funding decisions, and Section 8 discusses limitations before the conclusions.

2 Relation to technology roadmapping

The Technology Roadmapping Strategy (TRIS), developed at Politecnico di Torino over the past decade, provides a structured framework for technology prioritisation and developing maturation roadmaps [3, 4]. Recent developments have extended the method to represent the time and budget requirements for transitions between Technology Readiness Levels (TRLs) and to support roadmap selection under resource constraints [3].

The proposed ecosystem-readiness profile complements TRIS by addressing the implementability of a selected development path. Whereas TRIS supports identifying, prioritising, and planning technology maturation activities, the profile assesses whether the industrial capabilities, infrastructure, suppliers, and integration capacity needed to execute those activities are available within the relevant programme context.

Technology maturity describes what a technology has demonstrated. Ecosystem readiness concerns whether the programme can obtain the equipment, facilities, and expertise needed to implement it. Previous laboratory tests or flight experience provide evidence of maturity, but do not establish that suppliers can deliver the required configuration or that suitable testing and integration resources are available.

The model therefore treats qualification of the selected configuration as a separate requirement. The profile records evidence for the supply, testing and integration capabilities needed to implement it, alongside the work required where that evidence is incomplete.

This distinction draws on studies of innovation-system functions [5], organisational and activity interdependence [6], and the influence of ecosystems on technology substitution [7]. These studies explain how technology development and adoption depend on the organisations, resources, and supporting activities around a technology. For ECLSS, this means considering whether suppliers can deliver the required equipment, facilities can support its testing, and an integration team can assemble and verify the system.

The assessment follows the architecture from engineering performance through cost and maturation planning to the capabilities needed to execute that plan. This supports an additional question for HUMANS: which ground-development activities could establish evidence, interfaces and expertise useful to a later flight programme? Such transfer requires an explicit assessment of the receiving configuration and environment [8].

3 Case studies

HUMANS provides both the starting point and the primary motivation for this study. The terrestrial analog is intended to support research on the closure of the air and water cycles; accordingly, the technologies considered here belong to these two key life-support subsystems. The analysis examines how different configurations of air-revitalisation and water-recovery equipment can provide the required functions, what resources they demand, and which implementation capabilities they depend on.

A lunar lander is included as a secondary, exploratory application. This case is motivated both by the authors'

experience with lunar-lander systems and by the opportunity to examine whether the same analysis can be transferred from a terrestrial analog to a representative space-exploration context.

The analysis focuses on air revitalisation and water recovery. The selected functions include carbon dioxide removal, oxygen generation, carbon dioxide reduction, and water recovery from hygiene wastewater and humidity condensate. These functions provide a common basis for assessing how regeneration assumptions influence consumable demand, subsystem dependencies, and the development work required to implement each architecture. The scope reflects the distinct objectives of the two cases. For a short lunar sortie, the reduction in carried consumables may not justify the mass, complexity, qualification effort, and supporting infrastructure associated with regenerative processing. The analysis therefore compares stored-resource and regenerative alternatives. For HUMANS, regeneration has an additional experimental value, even where it is not strictly required to sustain a short-duration ground campaign. This study excludes solid-waste resource recovery, urine and brine processing, and food production. This exclusion defines the analytical boundary of the present study and does not imply that these functions lack value for longer-duration missions or future habitat research.

One HUMANS operational scenario is a fourteen-day occupied campaign with two crewmembers. As the longest envisioned crewed scenario, it defines the crew loads and duration used for resource sizing. The configurations considered must provide crew support while enabling experiments on air and water regeneration. The comparison is therefore between alternative life-support configurations that deliver the same required experimental capability, rather than simply between different means of accommodating two occupants. The purpose is to determine how the choice and organisation of air-revitalisation and water-recovery technologies affect the resources and implementation effort required by HUMANS. Project cost is the principal comparison objective, complemented by the availability of the equipment, facilities, and personnel required to assemble, integrate, and operate each configuration.

The lunar-lander case represents a short-duration sortie with four crewmembers, seven days on the lunar surface, and one additional day of active life-support operation during transit. Consumables are therefore sized for a total of eight days. The case is used as a reference design to investigate how the selected life-support functions affect a vehicle for which transported mass is a primary design consideration. Accordingly, the principal comparison objective is vehicle wet mass. In this context, regenerative technologies introduce a trade-off between reductions in stored consumables and the additional hardware, power, and integration resources required to provide regeneration. The lunar-lander application examines this balance over the selected mission duration. An open-loop configuration based on stored oxygen and water is included as a reference against which the mass consequences of introducing regenerative functions can be quantified.

Both cases consider the same core air- and water-processing functions. The regenerative air system combines carbon dioxide removal, oxygen generation through water electrolysis, and carbon dioxide reduction through the Sabatier reaction. Hydrogen produced by electrolysis is used in the Sabatier process to convert carbon dioxide into water, thereby returning part of the water required for oxygen generation. The architecture comparison considers an integrated arrangement following the functional organisation of the Advanced Closed Loop System (ACLS) [9], alongside an arrangement based on separate processing units. These alternatives allow the study to assess how different levels of functional integration affect the hardware, resource demands, and implementation capabilities required to provide the same life-support functions.

The water subsystem is represented by membrane-based treatment of hygiene wastewater and humidity condensate, informed at process level by ESA water-recovery development work [10]. Urine and brine processing are excluded from the selected comparison. The historical ultrafiltration/reverse-osmosis demonstration is used only to inform the assumed recovery performance and does not validate the specific membrane implementation adopted as the modelling proxy.

The two applications nevertheless generate substantially different water-processing loads because they assume different levels of hygiene provision. HUMANS uses a project-defined hygiene-water allowwater available for recovery and, consequently,nar lander uses a restricted allowance of 0.37 kg per person per day. Both cases include a separate flushing allowance of 0.30 kg per person per day. The larger hygiene-water stream in HUMANS increases the quantity of water available for recovery and consequently makes the external water requirement more sensitive to the assumed recovery performance.

The reference regenerative configurations use common literature-informed recovery assumptions. Carbon dioxide reprocessing is assumed to supply 40 % of the electrolysis feedwater requirement, following the ACLS functional design description [9]. Recovery of the eligible hygiene-water and humidity-condensate stream is set to 95 %, informed by ESA ground tests of ultrafiltration and reverse osmosis [10]. That test campaign included three 100-hour runs and reported water-recovery yields above 95 %, although the recovered water exceeded the potable-water limit for total organic carbon. The historical results therefore support the recovery assumption at process level but do not demonstrate either potable-water quality or the performance of the specific membrane proxy adopted here.

Both recovery percentages are shared modelling inputs for the two applications and regenerative architectures rather than demonstrated performance values for

each selected configuration. The water-recovery fraction applies only to collected, processable streams and does not itself establish potable-water quality. Its effect on system resources depends on the crew loads, mission duration, hygiene provision, and operating assumptions of each case.

The following section translates these case definitions into equipment, consumable, tankage, and power requirements. The resulting resource estimates provide the basis for assessing project cost and implementation feasibility in HUMANS, and vehicle wet mass in the lunar-lander application.

4 From architecture to an ecosystem-readiness profile

The assessment follows one architecture through four linked steps: resolve its equipment and resource demand; estimate its procurement and development costs; identify the maturation and implementation activities; and record the evidence needed to execute those activities.

4.1 Assessment logic

For each candidate architecture, the methodology first evaluates physical and operational performance. Resource balances quantify how water- and oxygen-recovery assumptions affect the amount of consumables that must be carried or supplied. The analysis then identifies the technology-development, qualification, testing, integration, and operational activities required to prepare each configuration for use. The assessment also examines the relationship between architecture choices and the supporting industrial ecosystem. It considers technology maturity, documented suppliers, potential suppliers with relevant adjacent experience, and dependencies on critical components, shared facilities, and specialist services. Where no supplier evidence or credible capability pathway is identified, the issue is recorded as an unresolved capability question rather than treated as evidence that the capability does not exist. This approach follows the principle that architecture development should assess the remaining development effort, possible maturation routes, and associated cost and schedule reserves. It also distinguishes demonstrated heritage from maturity achieved in the proposed configuration [8].

Evidence on industrial capability is often incomplete during early design. Each architecture is therefore assessed against a defined ecosystem and milestone. Supplier reviews, configuration tests, facility access and integration preparation are recorded as evidence needs or remaining activities. Their implications are examined through explicit work and resource scenarios. Architecture selection considers both physical performance and implementability. Unknown essential capabilities remain unresolved even when a scenario fits the available time. The profile identifies the conditions and evidence needed to support a recommendation, rather than inferring delivery probabilities from technology heritage or supplier records.

The methodology uses TRIS inputs and the curated database assembled for this research. These sources provide the engineering, technology-maturity, and industrial-capability evidence used to calculate resource requirements, identify development activities, and assess capability readiness.

4.2 Engineering and cost boundaries

Resource accounting separates installed equipment from tanks and initial supplies. For crew size n, active duration d, daily oxygen demand q_O and consumption margin s_f, the oxygen requirement is $M_O = s_f n d q_O$. The model uses $q_O = 0.84\,\mathrm{kg}$ per person per day and $s_f = 1.25$ as declared reference inputs. NASA's Baseline Values and Assumptions Document provides the broader resource-accounting basis [11].

Where electrolysis is installed, its make-up feedwater is $1.125 M_O(1 - r_A)$, with r_A the specified air-loop feedwater credit. The factor 1.125 follows from water-to-oxygen stoichiometry. The water ledger is

$$M_W = \max\{0, M_{\mathrm{use}} + M_{\mathrm{EVA}} + M_{W,O} - r_W M_{\mathrm{eligible}}\}, \tag{1}$$

where $M_{\mathrm{eligible}} = nd(q_H + q_C)$ contains hygiene wastewater and condensate. Urine and brine receive no recovery credit. Stored oxygen replaces electrolysis feedwater in the open reference. Shared processing hardware is counted once, even when it provides several functions. Equipment properties and crew-scaling rules remain generic study inputs requiring validation for selected units.

Lander ECLSS hardware and logistics enter the common iterative mass-estimating and stage rocket-equation model. The comparison reports the upper-retention wet-mass case, which retains the accounted initial consumable burden through ascent. Trajectory, power and thermal systems are modelled as explained in [12]. Equivalent system mass can also be calculated using NASA's accounting framework [13], but missing case maintenance hours and spare inventories prevent a complete ESM comparison here.

Published NASA and ESA estimates provide useful benchmarks for the calculated masses. ESA's Moon Village study assigns 850 kg to an ACLS sized for four crew, increasing to 935 kg with its 10 % mass margin [14]. The corresponding integrated air-loop contribution in the present lander model is 796 kg before margin and 995 kg after the adopted 25 % allowance. This agreement supports the estimated mass scale of the integrated air loop. At subsystem level, the historical NASA lunar-lander cost study specifies an ECLSS hardware mass of 1390 kg, compared with the present integrated equipment subtotal of 1312 kg [15]. At vehicle level, NASA's ESAS reference lander supports four crew for seven surface days at a total mass of 45.862 t, compared with 33.73–40.09 t for the candidates considered here [16]. These comparisons provide context for the

sizing results, although differences in equipment scope, propulsion and staging prevent direct validation of the complete architecture masses.

The comparison also identifies assumptions that require further substantiation. NASA's lunar-surface analysis assigns 467 kg to oxygen generation and 210 kg to Sabatier processing, including critical spares [17]. These estimates have a broader supportability scope than the historical component values used in the distributed candidate, whose spares inventory remains unresolved. Its calculated mass advantage therefore depends on reconciling equipment packaging, interfaces and spares between alternatives. Water savings likewise depend on which streams are collected: the cited NASA transit baseline includes an oral-hygiene water allowance but does not identify it as recoverable wastewater [18]. Excluding this assumed recovery stream increases the lander's regenerative water requirement from 122.2 to 133.5 kg, with all other inputs unchanged. The HUMANS equipment estimates additionally depend on the assumed crew-capacity scaling relationship. The results consequently support comparisons between explicitly defined sizing scenarios, while configuration-specific hardware and collection arrangements remain necessary to substantiate the final masses.

The budget results distinguish two programme boundaries. HUMANS has a project budget for design, services and hardware realisation, including the structure; operating costs are excluded. This is an available budget for the facility, not a separate cost estimate for each ECLSS architecture. The lander is represented by the HyCost [3, 19] lifecycle breakdown for the complete vehicle. NASA's Advanced Missions Cost Model (AMCM) provides the development-and-production relationship; launch, ground processing and operations are additional elements in HyCost's lifecycle accounting [3, 19]. NASA's cost literature likewise distinguishes development and production from launch and operations [20].

The figures retain each source's currency and cost boundary. A facility budget, a whole-vehicle model estimate and a commercial development-and-demonstration contract answer different questions. The contract benchmarks therefore provide financial-scale context rather than calibration targets or ECLSS unit prices. Architecture-specific costs require a documented allocation or a new model run with the relevant inputs.

4.3 Technology roadmaps and implementation work

TRIS supplies the maturation reference from the input technology level to a declared target [3]. A common TRL 9 endpoint is retained to show the reference trajectories consistently. That endpoint is not the commissioning requirement for HUMANS and does not replace verification of the selected lander configuration.

The TRL 9 references for air-processing technologies are grounded in documented operational heritage. Abney et al. identify the ISS-fidelity Sabatier system (ISAB) as TRL 9, distinguishing it from a developmental Sabatier implementation assessed at TRL 5 [21]. Similarly, Haggerty et al. state that Sabatier and water-electrolysis systems have achieved TRL 9 [22]. NASA's electrolysis state-of-the-art assessment assigns TRL 9 to the low-pressure ISS Oxygen Generation Assembly, while assigning lower TRLs to systems operating in other pressure regimes [23]. Nevertheless, TRL cannot be transferred directly between applications: the specific configuration, operating conditions, and system interfaces require case-by-case assessment [8]. These sources therefore provide a heritage-based benchmark for the present study, rather than a direct qualification of the proposed implementation.

The implementation roadmap encompasses requirements definition, procurement, configuration adaptation, relevant-environment demonstrations, verification, and integrated acceptance. Even technologies with flight heritage may require these activities when their operating conditions, configuration, or system interfaces differ from those of the heritage application [3, 8]. For each activity, the roadmap records predecessor relationships, duration ranges, resource requirements, and the status of supporting evidence. In this study, TRIS-based maturation estimates are combined with assumptions about procurement, integration, testing and resource availability to construct a scenario schedule. The supporting evidence is assessed separately to identify which assumptions require confirmation. The resulting completion dates are planning estimates, not confirmed delivery dates. The schedule accounts for shared resource constraints: activities using the same test facility or development bench are performed sequentially, while independent activities may run in parallel. It provides a feasible, priority-based sequence rather than an optimised programme schedule. Two schedule views are analysed. The scenario schedule applies the stated assumptions on activity durations, maturity, and resource availability to illustrate a possible execution path. The evidence-only schedule reports a completion date only when the supporting maturity assessment, duration estimate, and resource-access assumptions are adequately substantiated. Accordingly, scenario completion dates should be interpreted as analytical planning outputs, not as supplier commitments or evidence of demonstrated programme readiness

4.4 Definition of the readiness profile

This study develops an initial implementation of ERI as an ecosystem-readiness profile. The profile describes what is known about a selected architecture, what must still be done before it can be used, and which external resources are needed to complete that work. It connects the technical requirements of the architecture with the suppliers, facilities and organisations responsible for its delivery, integration and verification.

For architecture a, ecosystem e, and required milestone t_n, the profile is written as

$$\mathcal{R}(a, e, t_n) = \{\mathcal{E}, \mathcal{W}, \mathcal{D}; C_W, M_T, M_C\}. \quad (2)$$

Here, $\mathcal{E}$ records the requirements and their supporting evidence; $\mathcal{W}$ records the remaining work, including procurement, adaptation, integration and verification; and $\mathcal{D}$ records the dependencies between that work and the resources needed to complete it. For example, an integration activity may depend on both equipment delivery and the availability of a suitable structure.

The three accompanying figures of merit describe the estimated cost of the remaining work, C_W, the time available between its completion and the required milestone, M_T, and the remaining budget, M_C, where a budget covering the same work is available. These quantities help assess whether an implementation path is feasible, but they are reported separately rather than combined into a readiness score.

Each requirement retains one of three evidence states. *Supported at the stated scope* means that the available evidence supports fulfilment of the requirement under the conditions assessed. *Documented shortfall* means that the evidence identifies a requirement that is not yet met. *Unknown* means that the available information is insufficient to determine whether the requirement is met. An unknown is therefore an information gap, rather than evidence of either success or failure. The three states also call for different actions. This study names the resulting mapping a *gap-to-action typology* (Table 1). An unknown is resolved by an information action, such as a request for information, a market consultation or a supplier audit, which changes what the programme knows without changing the design or the supply base. A documented shortfall requires an investment action, such as a qualification campaign or a new supply route. A dependency on the structure, a test facility or a shared resource requires an infrastructure action, such as an access agreement, an additional bench or protection of the gating milestone. Keeping the three classes apart prevents a programme from funding development where a search would suffice, or from searching where the capability is known to be absent.

In the current HUMANS profile, every essential requirement of both candidates is in the unknown state and none is a documented shortfall. The European process routes for the air-processing and membrane water units are evidenced, but delivery of the configured units is unconfirmed; humidity control has research or catalogue leads only; and atmosphere monitoring, trace-contaminant control and water polishing depend on a similarity assessment of industrial equipment. The typology therefore places information actions first in the HUMANS sequence, ahead of development funding. The infrastructure class holds the two dependencies that shape the HUMANS schedule: the structure milestone and the development bench shared by the air and water chains (Section 5.5).

Evidence must address the configuration and conditions being assessed. A successful process demonstration can establish relevant experience without confirming that the selected equipment will meet the requirements of the proposed architecture. Similarly, identifying a suitable laboratory does not confirm access to it during the required test period, and identifying an organisation with integration expertise does not establish that personnel have been assigned. The profile records these distinctions so that demonstrated capabilities and outstanding conditions remain visible.

Table 2 summarises the questions addressed by the profile.

For a defined scope of remaining work, the direct cost is

$$C_W = \sum_{j \in \mathcal{W}} c_j, \tag{3}$$

where c_j is the estimated cost of work package j, counted once. Calculating C_W requires estimates for the activities actually included in the profile. The available programme-level financial records provide useful cost context, but they do not allocate costs to all these activities. They therefore cannot be treated directly as the cost of the remaining architecture-specific work.

Schedule margin compares the estimated completion date with the required milestone:

$$M_T = t_n - t_{\mathrm{finish}}. \tag{4}$$

The required milestone is the date by which the capability must be available with its required verification complete. A positive margin means that the evaluated scenario finishes before that date; a negative margin means that it finishes afterwards.

For HUMANS, the specified target is Q1 2029. Because this target is a quarter rather than an exact date, and activity durations are represented by ranges, the schedule margin is also reported as a range. The lower bound compares the latest estimated completion with the beginning of Q1, while the upper bound compares the earliest estimated completion with the end of Q1. This expresses the range of margins implied by the stated planning assumptions.

A positive margin shows that the assumed activities could fit within the available time. It does not confirm that equipment will arrive as assumed, that facilities will be accessible, or that testing will achieve the required results. The evidence supporting these assumptions is therefore assessed separately and retained in the profile.

Cost margin is defined only when an approved budget B_W covers the same activities as C_W:

$$M_C = B_W - C_W. \tag{5}$$

The HUMANS design-and-hardware budget covers a broader scope than the selected ECLSS activities. Without a corresponding allocation to those activities, subtracting their estimated costs from the full programme budget would give a misleading margin. Architecture-specific cost margin therefore remains unresolved in this study.

This treatment follows NASA's distinction between measurable decision criteria and mandatory requirements [24]. Cost and schedule inform the comparison

Table 1: Gap-to-action typology: the class of programme action that each unresolved condition in the readiness profile calls for, with the current HUMANS assessment.

Profile condition	Action class	Typical instruments	Who acts	HUMANS, current profile
Unknown requirement ($\mathcal{E}$)	Information	Request for information, market consultation, supplier audit, similarity assessment of industrial equipment	Programme office with suppliers	All essential requirements of both candidates: configured air and water units unconfirmed; humidity control with research or catalogue leads only; monitoring, trace-contaminant control and water polishing awaiting a similarity assessment
Documented shortfall ($\mathcal{E}$) or remaining development ($\mathcal{W}$)	Investment	Breadboard and relevant-environment work, qualification campaign, supplier development, new supply route	Programme with its funding body	No documented shortfall; the scenario development work follows from the assumed maturity levels (Fig. 5)
Dependency on structure, facility or shared resource ($\mathcal{D}$)	Infrastructure	Access agreement, shared test infrastructure, additional bench, protection of a gating milestone	Host institution or facility owner	Structure available by end 2027 gates all equipment work; one development bench is shared by the air and water chains (Section 5.5)

of implementation options, while essential requirements must still be satisfied. A favourable cost or schedule estimate cannot compensate for a documented shortfall or establish that an unknown requirement has been met.

4.5 Using the profile for sequential decisions under uncertainty

The profile also provides a basis for identifying what to do next. As work progresses, tests, supplier information and integration activities produce new evidence. This evidence can change both the assessment of the architecture and the work required to complete it. The decision problem is therefore sequential: the most appropriate next action can depend on the outcome of an earlier action.

A partially observable Markov decision process (POMDP) provides a framework for describing such decisions when some relevant conditions are uncertain [25]. In this study, the profile is the engineering record, while the POMDP describes how that record could be used to compare actions and update decisions as evidence becomes available. The POMDP does not redefine ERI as a probability or a single readiness score.

For HUMANS, an unresolved equipment interface illustrates this relationship. The profile would record the interface requirement, the missing evidence and the integration activities that depend on it. A possible next action could be to obtain additional documentation or perform a configuration test. A satisfactory result could support proceeding with integration, whereas an unsatisfactory result could identify a need for modification and further testing. Each outcome would update the evidence record and could change the remaining work, its cost and its completion date.

To express this decision process mathematically, let x represent the joint condition of supply, configuration performance, interfaces and test access. At decision step k, the belief $b_k(x)$ is a probability distribution over these possible conditions, based on the evidence available at that step. It represents uncertainty about the conditions, rather than the readiness of the programme as a whole. The observed date and remaining budget are also available to the decision model. The qualitative evidence states in the profile do not, by themselves, provide the numerical probabilities required for this belief.

After an action u_k produces an observation o_{k+1}, the belief is updated as

$$b_{k+1}(x') = \eta\, O(o_{k+1} \mid x', u_k) \sum_{x} T(x' \mid x, u_k) b_k(x). \tag{6}$$

Here, T describes the probability that the action changes the condition from x to x', while O describes the probability of obtaining the observation under the resulting condition. The factor η ensures that the updated probabilities sum to one. The update therefore accounts for both the possible effect of an action and the information provided by its observed result.

In programme terms, these actions correspond to existing instruments: documentation reviews and requests for information change only the knowledge of the state; configuration tests and qualification campaigns both inform and change it; facility-access and supplier-development agreements change the ecosystem itself. The three action classes of Table 1 therefore also define the action set of the decision model: information actions change the belief through O only, whereas in-

Table 2: Contents of an architecture-specific readiness profile.

Element	Questions addressed	Typical programme action
Required functions	Does the selected architecture provide the required functions, and what evidence supports this assessment?	Technology maturation through the TRIS roadmap.
Supply	Are suitable suppliers or development routes identified? Is delivery of the selected configuration confirmed?	Issue an RFI or conduct a market consultation where the supply base is unknown; fund supplier development or qualification where a shortfall exists.
Facilities	Which tests are required, and are suitable facilities accessible when needed?	Establish an access agreement or use shared test infrastructure.
Integration	Are interfaces and verification activities defined? Are the necessary personnel assigned?	Provide staffing, training, and measures to preserve team continuity.
Remaining work	What must still be purchased, developed, adapted, demonstrated or verified?	Fund the remaining work through defined work packages.
Dependencies	Which deliveries, facilities or shared services could prevent completion?	Secure a second source or implement an action to remove the dependency.
Figures of merit	What would the remaining work cost, and how much time and budget would remain?	Reallocate budget and/or schedule margin as required.

vestment and infrastructure actions can also change the state through T.

Different actions affect the assessment in different ways. Reviewing existing documentation may improve knowledge without changing the equipment, whereas modifying equipment can change its performance and create a need for further verification. Dependencies must also be retained: architectures that share a supplier, facility or interface may be affected by the same uncertainty. Independent probabilities for these conditions are therefore not assumed.

The actions considered for HUMANS must follow the practical sequence of the work. Equipment must be procured before installation, the structure must be available before integration, and tests require suitable resources and sufficient time. A numerical decision model could compare feasible actions using an agreed programme objective, such as minimising expected additional development and testing cost while meeting a specified completion requirement and an agreed limit on the risk of delay. These choices would need to be defined before calculating a policy, meaning a rule for selecting the next action from the evidence available. The objective would guide action selection without replacing mandatory verification and acceptance requirements.

This first implementation establishes the readiness profile and dependency structure and formulates how they could support sequential decisions. It does not yet calculate a calibrated POMDP policy for HUMANS. Doing so would require initial probability estimates, configuration-specific action costs and durations, probabilities of action outcomes, and estimates of test reliability. Earlier exploratory calculations remain illustrations of their stated assumptions. The profile can nevertheless be used now to identify unresolved requirements, the activities affected by them and the evidence needed to support the next decision.

5 Architecture, cost and implementation results

This section answers three questions in turn: what each architecture requires physically, what the available budgets cover, and what work, evidence and dependencies stand between each architecture and its milestone. Each subsection includes its reading for programme decisions.

5.1 The physical trade changes with the case

Figure 1 separates the equipment and supply contributions. In the lander reference, regeneration reduces carried water from 179.8 to 122.2 kg, but the equipment burden dominates the short mission. The open, integrated and distributed candidates have accounted ECLSS masses of approximately 614, 1530 and 953 kg, respectively. Their upper-retention vehicle wet masses are 33.73, 40.09 and 35.80 t. These are conditional outputs of the reduced sizing model, not reconstructed flight-vehicle masses.

The HUMANS integrated and distributed candidates have accounted masses of approximately 1127 and 763 kg. Their required research functions are retained in both alternatives. The lighter candidate does not automatically provide the easier procurement or integration path, and physical mass is not the terrestrial project's principal selection criterion.

The hygiene-water assumption explains an additional difference between the cases. A ten-percentage-point reduction in eligible-stream recovery adds 8.45 kg of water to the lander and 27.83 kg to HUMANS. The same reduction in air-loop feedwater credit adds 3.78 and 3.31 kg, respectively. These results identify the water stream that needs particular attention in a ground demonstration; they do not establish the performance of an unselected processing unit.

Programme reading. For the short sortie, regeneration adds more equipment mass than it saves in consumables, so the case for developing it rests on later, longer

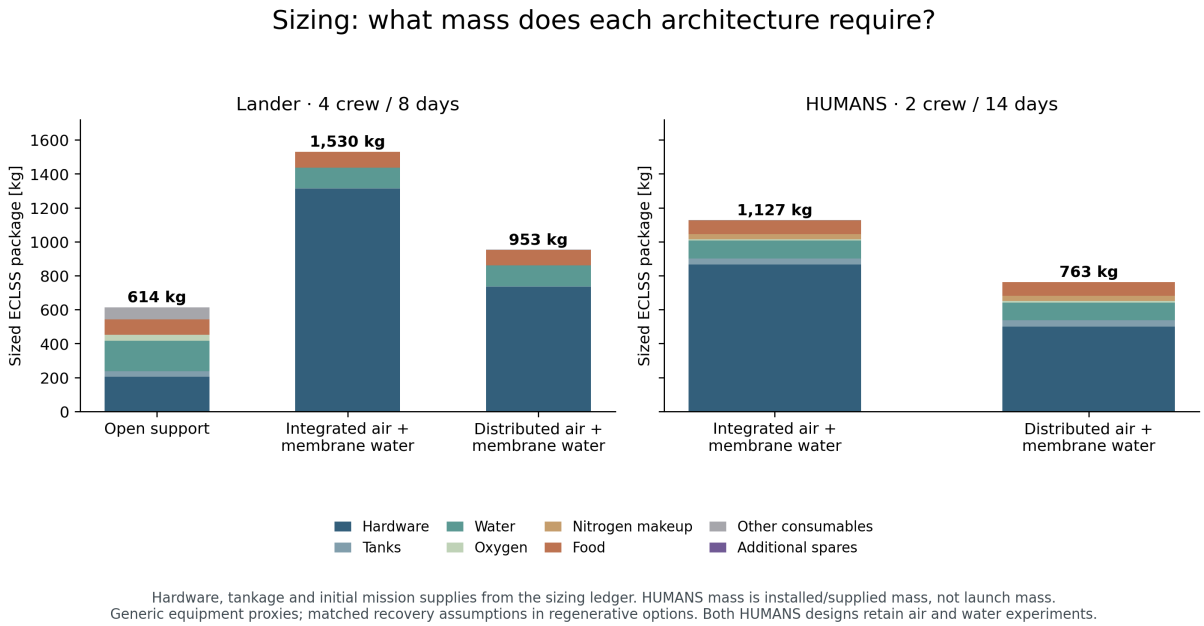


Fig. 1: Architecture resource breakdown from the common sizing ledger. The lander uses four crew and eight active life-support days; HUMANS uses two crew and fourteen days. Regenerative candidates share the specified recovery assumptions. HUMANS quantities are installed equipment and initial supplies. Additional case-specific spares remain unresolved.

missions rather than on this vehicle. Of the two recovery assumptions, the water-recovery fraction has the larger effect on HUMANS supplies, which makes it the first quantity a ground demonstration should measure.

5.2 *Programme budgets and lifecycle costs have different boundaries*

Figure 2 presents the financial basis for each case. The HUMANS project budget is around 1.320 million euros: around 22% for design and services and 78% for hardware and realisation. It includes the facility structure and equipment, while operations are excluded. This project-provided budget is grounded in the facility being developed; it is not a separate procurement quotation for either ECLSS candidate.

The HyCost HLS run reports a whole-lander lifecycle estimate of approximately 2.67 billion US dollars. Its breakdown includes development and testing, production, launch, ground processing and integration, life-support operations and crew mission operations. Development and production together account for approximately 1.95 billion dollars; the remaining elements are included in the reported lifecycle total. These are whole-vehicle amounts, not the development cost of the selected ECLSS alone.

Published human-lander awards provide an external check on the financial scale (Fig. 3). NASA announced an original 2.89-billion-dollar award to SpaceX in 2021 and a 3.4-billion-dollar award to Blue Origin in 2023 [26, 27]. These awards cover lander development and demonstration, including uncrewed and crewed demonstrations [27, 28]. The HyCost result is on the same broad billion-dollar scale, supporting an order-of-magnitude reasonableness check.

For the architecture trade, a separate HyCost sensitivity gives relative whole-vehicle development-cost indices of 100 for the open reference, 118.6 for the integrated configuration and 107.1 for the distributed configuration, with other cost drivers fixed. These indices describe the modelled consequence of changing the architecture's mass; they neither allocate the archived lifecycle total to ECLSS nor establish absolute configuration-specific costs.

Programme reading. Neither total can serve as the budget against which the architecture-specific work is judged. A cost margin for HUMANS therefore needs an allocation of the project budget to the ECLSS work packages, which is a project decision rather than a modelling gap.

5.3 *TRIS and the capability roadmap answer different questions*

Figure 4 presents the TRIS reference paths for the selected technology proxies. The membrane, thermal-amine and integrated-air inputs require approximately 7.16, 4.71 and 2.33 reference years to the common TRL 9 endpoint. These durations follow the shared TRIS profile and generic input levels; they are not configuration-specific development forecasts.

The capability roadmap in Fig. 5 shows the work omitted by a maturity-only view: equipment procurement, adaptation, shared test resources and integrated acceptance. It uses case-specific work assumptions rather than the common TRL 9 endpoint. The dis-

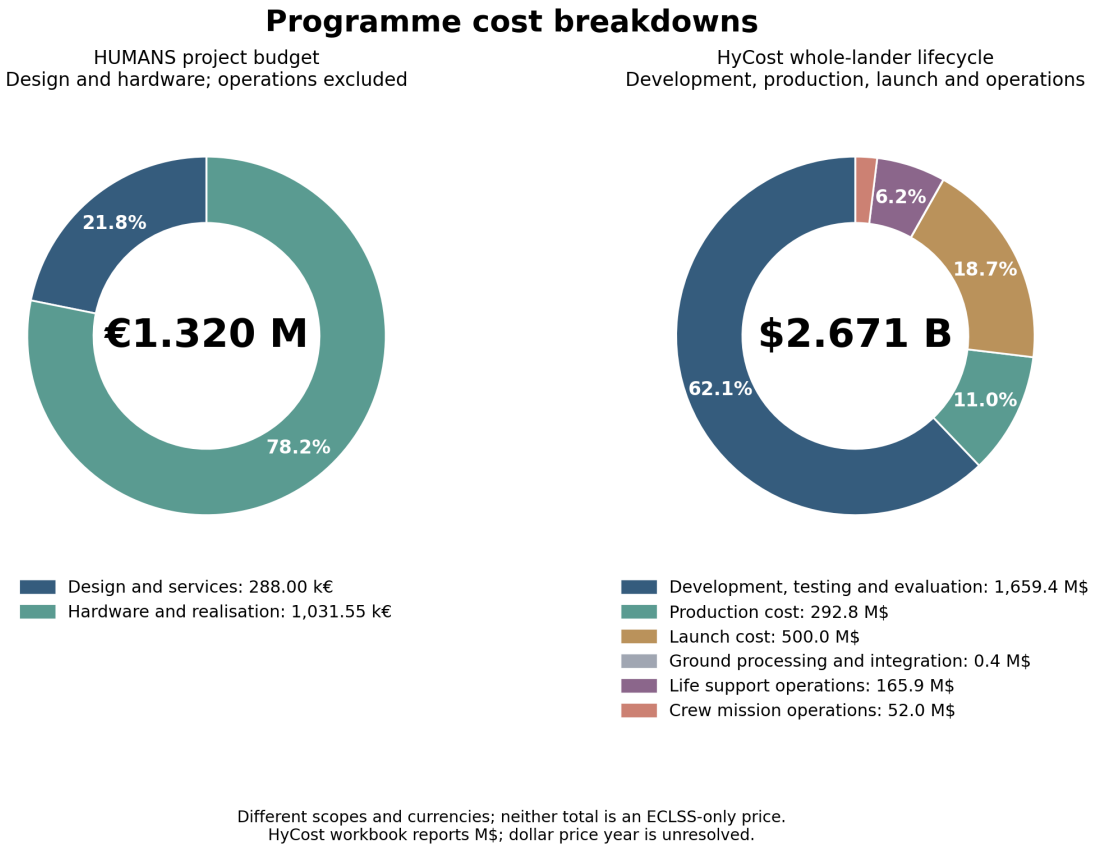


Fig. 2: Programme financial breakdowns with their respective scopes and currencies. HUMANS shows the project design-and-hardware budget, excluding operations. HyCost shows the archived whole-lander lifecycle estimate, including launch and operations. Development excludes production in its plotted component, so production is counted once. The totals are not ECLSS-only prices.

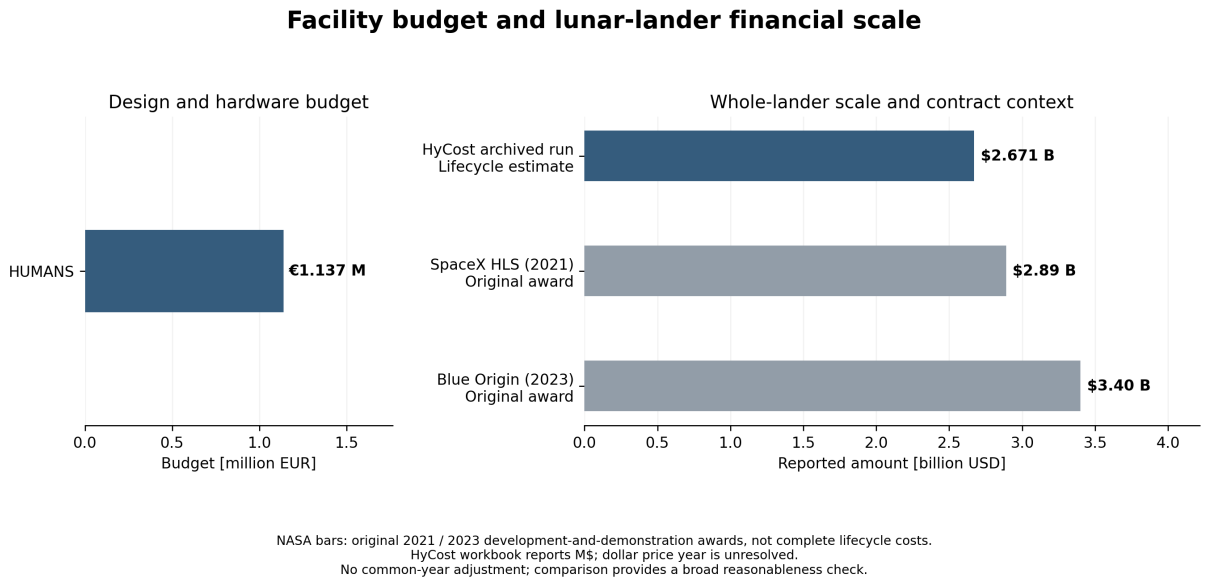


Fig. 3: HUMANS project budget and whole-lander financial context, using separate monetary scales. NASA bars show the original 2021 SpaceX and 2023 Blue Origin development-and-demonstration awards, not current contract totals or complete lifecycle costs [26, 27]. The HyCost estimate includes operations.

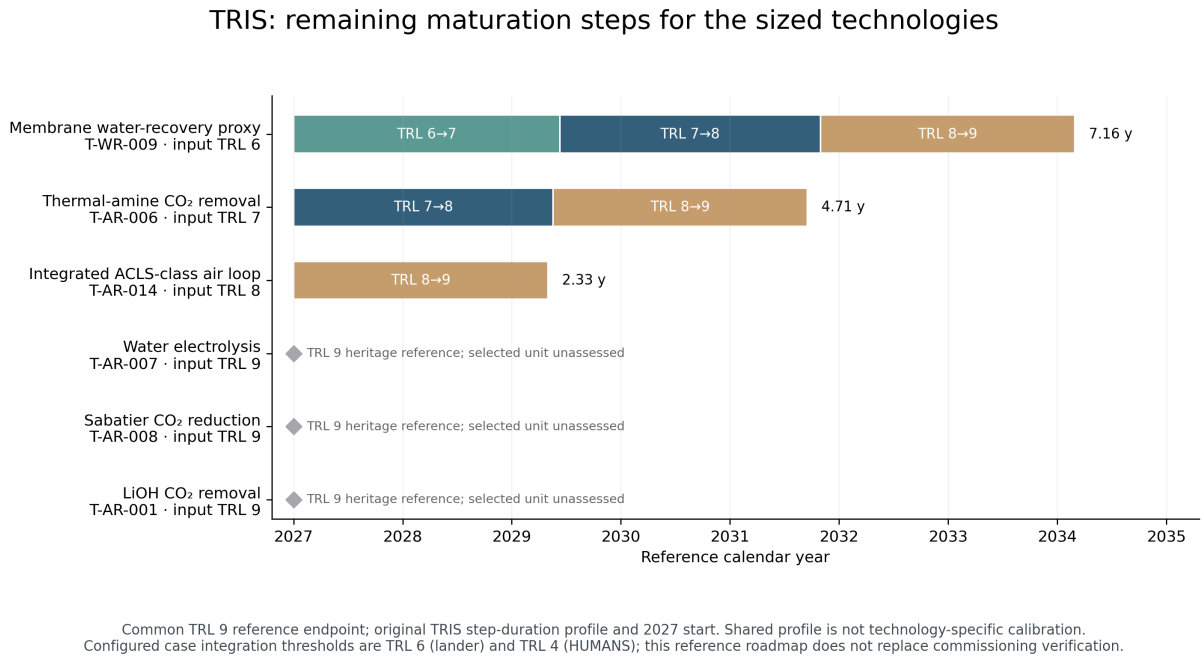


Fig. 4: TRIS reference maturation paths for the air- and water-technology proxies used in the architecture comparison. Sabatier and electrolysis TRL 9 inputs refer to documented heritage systems [21, 22, 23]. The common endpoint does not define HUMANS commissioning, establish a lander delivery date or qualify the selected equipment.

tinction prevents a generic technology level from being interpreted as immediate availability of the required subsystem. *Programme reading.* For a funding body, this is the main difference between the two views. A maturity-gated criterion would treat the TRL 9 electrolysis and Sabatier units as available, whereas the capability roadmap shows that their procurement, adaptation and verification still occupy the implementation path.

Table 3 reports the elapsed implementation time and schedule-margin envelope. Both lander ranges cross their illustrative deadline. With structure availability included, the HUMANS integrated scenario finishes between 13 May 2028 and 20 February 2029; the distributed scenario finishes between 22 June 2028 and 20 June 2029. The longer integrated endpoint falls within Q1, while the longer distributed endpoint extends beyond it. These conditional calculations do not demonstrate that the project can meet its target: task estimates, delivery commitments, test completeness, facility access and staffing remain unresolved.

These outcomes explain the value of retaining a profile. The integrated candidate completes the declared implementation sequence sooner in the present scenarios, while the distributed candidate has lower accounted physical mass. Neither observation resolves the unknown procurement prices, required verification or delivery capability. The next decision concerns which configuration evidence and work estimates are needed to substantiate the trade.

A historical NASA contractor study provides an external benchmark: its lunar-lander crew-module ECLSS had a cost-model hardware mass of 1390 kg, an estimated 41-month development period and a development cost of 118.522 million 1988 US dollars, excluding production [15]. NASA's New Start Inflation Index converts this to approximately 355 million FY2026 US dollars; the FY2026 index value is projected [29]. The present integrated and distributed hardware masses, 1312 and 735 kg, are within a factor of two of that historical value, although equipment boundaries differ. The 3.4-year historical development estimate is also on the same broad multi-year scale as the TRIS reference paths, whose starting maturity and scope differ. This historical ECLSS-only development estimate complements the whole-lander financial context. It does not represent a share inferred from the HyCost total or establish the cost of full flight qualification for the present ECLSS configurations.

5.4 Evidence gaps and dependencies remain visible

ESA's ACLS report supports European integrated air-process heritage [9]. ESA's historical water-recovery work supports process and testing experience [10]. These sources do not confirm supply of the selected equipment, and the historical water process is not the same as the membrane proxy. A relevant monitoring method provides another identifiable evidence contribution [30]. Its existence does not settle water acceptance for the proposed flight configuration.

The profile distinguishes those contributions from the unresolved questions: who can deliver the configured unit, at what price and date; which facility can verify

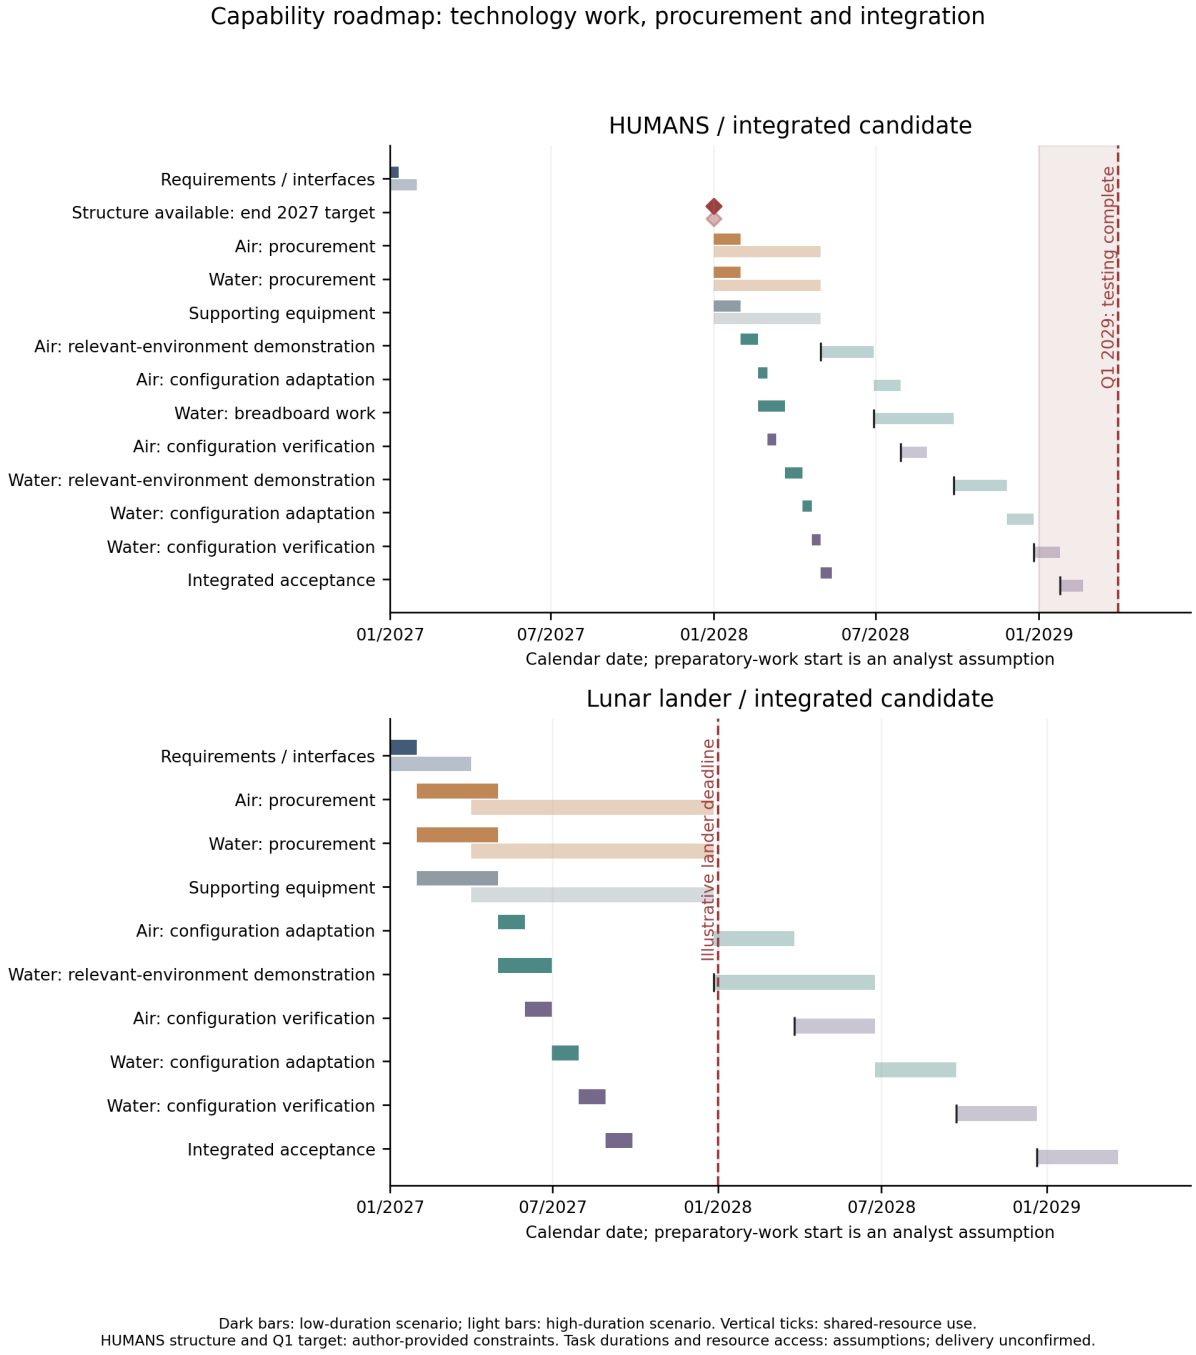


Fig. 5: Implementation work for the integrated candidates. Dark and light bars show the shorter and longer task-duration scenarios. HUMANS procurement waits for the end-2027 structure target; the shaded Q1-2029 window denotes readiness with integration/testing complete. The lander deadline remains illustrative. In the terms of Table 1, development and verification bars correspond to investment actions, and the structure milestone and shared-resource ticks to infrastructure dependencies; information actions do not yet appear as scheduled activities.

it; and which team can integrate and operate it. A route absent from the reviewed evidence is a search gap, not proof that no supplier exists. Supplier records also do not establish independent alternative production routes. Each of these questions has an owner: suppliers answer supply and price through a request for information or quotation; test-centre agreements cover facility access; and the programme staffing plan covers integration capability. Recording the owner alongside the evidence state turns the list of unresolved items into a list of actions.

The structure-availability milestone is a programme prerequisite for the HUMANS equipment sequence. Separately, the functional dependency model identifies shared air-processing hardware and services whose loss removes required functions. These are different uses of

Table 3: Conditional implementation scenarios. Elapsed time starts on the assumed 1 January 2027 preparation date and includes waiting for the HUMANS structure. HUMANS margin bounds combine the beginning/end of Q1 with the longer/shorter completion endpoints; lander margins use one illustrative deadline. Ranges are not confidence intervals. Cost margin and evidence-supported delivery dates remain unresolved.

Case / configuration	Elapsed time (days)	Direct work (kEUR)	Schedule margin (days)
Lander / integrated	270–810	115–345	-445 to $+95$
Lander / distributed	300–900	175–525	-535 to $+65$
HUMANS / integrated	498–781	165–495	-50 to $+322$
HUMANS / distributed	538–901	265–795	-170 to $+282$

dependency: the structure gates development, whereas the functional graph describes consequences of equipment or service loss. Procurement alternatives do not constitute installed redundancy. Both dependency views inform the evidence and tests needed before acceptance.

Programme reading. The HUMANS readiness date is governed by infrastructure availability rather than by technology maturity, so schedule protection comes from securing the structure, not from accelerating technology development. This reading follows from the project plan, which identifies the structure as the blocking item, rather than from the model; the model shows what happens to the equipment work once that milestone is fixed.

5.5 Schedule consequences of the declared stresses

Figure 6 shows the schedule-margin range for the reference scenario, each individual stress and the combined stresses. It retains the structure prerequisite for HUMANS. These scenarios perturb equipment work after the structure target; they do not assess uncertainty in the structure completion date. A sensitivity ranking within this task model cannot overturn the project's identification of the structure as its blocking item.

Some additional activities fit within existing schedule slack, while others extend the controlling chain. The figure identifies timing consequences without treating the exploratory work-package amounts as either the HUMANS budget or the lander lifecycle cost. A monetary estimate of the effect of each stress requires a supported change in the relevant programme cost model. Under all declared stresses the lander margin against its illustrative deadline becomes negative across the whole range, while HUMANS retains a positive upper bound.

The comparison that does not depend on that deadline is the change in margin that each stress produces within each case. The supplier-delay stress adds the same 60 to 180 days to air-equipment procurement in both cases. In HUMANS, it reduces the margin by the full 60 to 180 days: the scenario schedule gives the shared development bench to the air unit first, so a late air delivery also delays the water chain, which waits for the bench. In the lander, the water chain controls the schedule, and the same delay reduces the margin by 30 to 90 days. The additional-development stress acts the other way: it reduces the lander margin by 150 to 420 days and the HUMANS margin by 30 to 90 days. Part of this second difference comes from the scenario maturity levels. The stress sets the air chain to the same level in both cases, which removes two assumed maturity steps in the lander and one in HUMANS. These durations and levels are analyst scenarios, so the results show mechanisms rather than forecasts.

Programme reading. In the ground facility, the schedule is most exposed to equipment delivery and to access to the shared bench, which supplier engagement, bench planning or an additional bench can address. In the lander, it is most exposed to the maturity assumed at the start: a plan that assumes more maturity has more to lose if that assumption fails. The same technology proxies therefore carry different programmatic exposure in the two development contexts.

6 HUMANS as a public research infrastructure for capability development

HUMANS can be understood not only as an ECLSS development facility, but also as a public research infrastructure that builds capabilities for subsequent flight programmes. Its value therefore extends beyond the facility itself to the technical evidence, integration experience, specialist skills, test procedures, and operational records generated during its design, integration, testing, and operation.

Viewed through the innovation-system functions identified by Hekkert [5], HUMANS could contribute to several forms of capability development. It could support knowledge development by generating process data under realistic loads and operating conditions. It could enable knowledge diffusion through records, procedures, and lessons learned that are usable by later programmes. It could provide experimentation capacity by acting as a test bed through which suppliers without flight heritage can develop and assess relevant technologies. Finally, it could contribute to resource mobilisation through the training of personnel and the maturation of test methods, facilities, and operating procedures.

NASA's use of the ISS as a testbed for exploration ECLSS provides a precedent for linking operational experience with subsequent system development [31]. HUMANS would generate evidence in a terrestrial environment, however, and the relevance of that evidence to a flight application would need to be assessed against the requirements, configuration, and operating environment of the receiving system. NASA's technology-assessment guidance similarly requires heritage to be reassessed

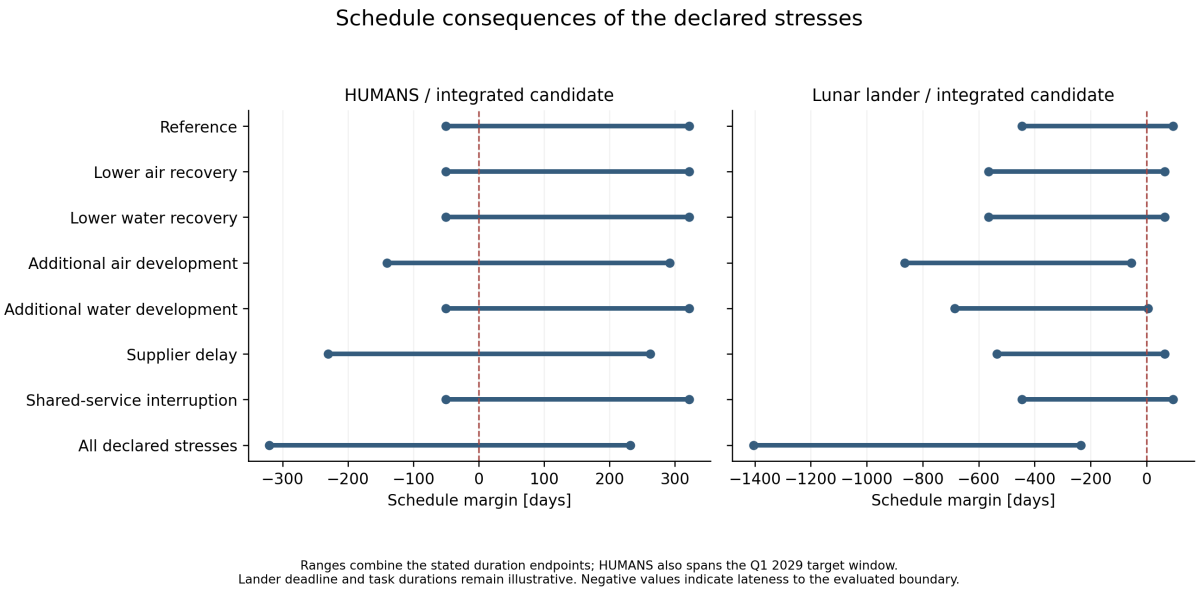


Fig. 6: Schedule-margin scenarios for the integrated candidates. The ranges combine the declared task-duration endpoints; HUMANS also spans the beginning and end of its Q1-2029 target window. The lander deadline remains illustrative. These scenarios have no assigned probabilities and do not establish delivery commitments. Programme cost totals are reported separately in Figs. 2 and 3.

when either the configuration or the operating environment changes [8].

Accordingly, the capability-development activities associated with HUMANS should be interpreted as a proposed pathway for reducing uncertainty and building competence, rather than as completed demonstrations or an automatic increase in flight TRL. Architecture selection should therefore consider both the resources required to implement HUMANS and the longer-term capabilities, evidence base, and supplier ecosystem that its development could make available to future programmes.

6.1 Selecting tests for future engineering questions

A useful starting point is to identify a question that a future flight programme must answer and then determine which elements of that question HUMANS could investigate. For example, a programme considering membrane-based water processing would need to understand water recovery, product-water quality, fouling behaviour, and maintenance requirements. Testing the selected processing chain in HUMANS could characterise these quantities under defined operating conditions and identify the modifications or additional experiments required. Similarly, air-processing tests could examine power demand, recovery performance, control stability, and interactions between connected items of equipment.

The connection to ERI is direct: the profile identifies an unresolved requirement, a HUMANS activity investigates it, and the resulting evidence updates the assessment. A successful ground test may support satisfaction of a requirement within the tested conditions, while leaving compliance under flight conditions unresolved. Conversely, an unsuccessful test may identify a documented shortfall and define the corrective work required. In either case, the profile becomes more informative by clarifying the work that remains.

To be useful beyond HUMANS, the resulting records should identify the equipment configuration, operating loads, measurement methods, acceptance criteria, and observed results. Maintenance interventions and responses to off-nominal events should also be documented. A future design team can then determine which findings are applicable to its architecture and which require further testing. Table 4 illustrates this relationship.

6.2 Implications for architecture selection

The HUMANS architecture should therefore be evaluated against the experiments and development activities that it is intended to support. Relevant considerations include whether individual components can be accessed, operating conditions adjusted, interfaces modified, and the required measurements acquired. Such features may affect the facility's value for capability development even when they do not minimise equipment mass or purchase cost.

For example, an integrated package and a set of separate process units may provide different experimental opportunities. However, the current data do not establish which candidate configuration would provide superior access for testing, modification, or diagnosis. Resolving this question requires engagement with the intended operators and suppliers.

For a subsequent flight programme, the evidence generated through HUMANS could inform requirement development, equipment selection, supplier engagement,

Table 4: Proposed HUMANS activities, their innovation-system functions, and their potential contribution to a subsequent flight programme. These activities describe future work rather than demonstrated HUMANS results.

Capability	Innovation-system function	What HUMANS could establish	What a later flight programme must assess
Air processing	Knowledge development	Recovery, power demand, load response, and interactions between connected units	Applicability to the selected flight configuration, mission loads, and operating environment
Water processing	Knowledge development	Recovery, product-water quality, fouling behaviour, and maintenance under recorded feed conditions	Applicability to flight wastewater streams, operating conditions, and acceptance requirements
Integration	Knowledge diffusion	Defined interfaces, integration procedures, and observed responses to faults	Configuration changes, interface updates, and remaining verification for the flight system
Operations and people	Resource mobilisation	Maintenance experience, task durations, training needs, and interruption-recovery procedures	Suitability for crew autonomy, available resources, and flight operating constraints
Test infrastructure	Resource mobilisation; experimentation	Facility operating limits, measurement capability, and experience conducting relevant tests	Which development tests the facility can support and which require other facilities

and the definition of remaining development work. The flight programme would nevertheless need to demonstrate compliance with its own environmental, performance, reliability, and integration requirements. No quantitative reduction in flight-development cost or duration is assigned in this study, because neither the receiving configuration nor the work that HUMANS evidence could credibly eliminate has yet been defined.

The two numerical cases considered in this study represent separate implementation scenarios. Because the illustrative lander deadline precedes the HUMANS Q1–2029 readiness target, HUMANS cannot be presented as supporting that particular schedule. Instead, the proposed pathway concerns a later flight programme. This initial implementation of ERI records the relevant engineering question, the HUMANS activity that could investigate it, the evidence that should be retained, and the verification that would still be required for flight.

7 Implications for programme and funding decisions

Table 5 collects the implications of the results for programme selection and funding. Each implication is labelled by the strength of its basis in this study: a *model result* is computed here under declared assumptions; a *case observation* comes from HUMANS project information rather than from the model; a *proposal* has not yet been tested.

8 Discussion and limitations

The engineering comparison and readiness profile support complementary decision-making processes. The sizing analysis identifies the physical implications of each architectural option, while the cost breakdowns distinguish between the HUMANS design-and-hardware budget and the HyCost whole-lander lifecycle estimate, which also includes operations. TRIS provides a reference for technology maturation, whereas the capability roadmap and evidence ledger identify the activities, resources, and evidence required to implement the selected configuration within a specific industrial ecosystem. These elements are linked through common functions, equipment allocations, and work packages.

The readiness profile is intentionally structured to preserve the reason for an unfavourable or incomplete assessment. A missing supplier commitment, an undefined verification campaign, and an exceeded schedule margin, for example, each require a different response from programme managers, suppliers, or facility operators. Combining such conditions into a single aggregate score would obscure these distinctions and reduce the usefulness of the assessment. This follows guidance on composite indicators for policy use, which cautions that aggregation can hide the reasons behind a score and allow strong dimensions to compensate for weak ones [32]. Retaining native cost and time units also enables readers to examine the underlying activity estimates without first accepting an additional readiness scale.

The numerical results remain conditional on the available evidence. Generic equipment properties, incomplete spare inventories, and a reduced-order lander model constrain the accuracy of the current estimates. Each limitation, however, corresponds to a practical next action: supplier quotations and configuration confirmation can reduce the cost margin; a verification and facility-access plan can resolve test dependencies; interface-definition work can close remaining integration questions; and configuration-specific HUMANS records can address the transfer questions relevant to a receiving flight programme. The readiness profile can then be updated as these actions generate new evidence.

Table 5: Implications for programme and funding decisions, with the basis and strength of each in this study. Model result: computed here under declared assumptions. Case observation: HUMANS project information, not a model output. Proposal: not yet tested.

Implication	Basis in this study	Strength	Decision and actor	Typical instrument
Maturity is not availability	TRL 9 air-processing units still carry procurement, adaptation and verification work in the capability roadmap (Fig. 5); heritage must be reassessed for a new configuration [8]	Model result, consistent with guidance	Selection and eligibility criteria; funding body	Request implementation evidence alongside TRL
Where delay enters depends on the development context	Supplier delay propagates fully through the shared bench in HUMANS; additional development dominates in the lander (Section 5.5)	Model result, conditional on scenario durations and maturity levels	Schedule protection; programme manager	Supplier engagement and bench planning for the ground facility; maturity verification before commitment for the flight vehicle
Infrastructure can bind before technology	The structure milestone gates all HUMANS equipment work (Fig. 5)	Case observation	Facility schedule; host institution	Protection of the structure milestone
Resolve unknowns before funding development	All essential HUMANS requirements are unknown and none is a documented shortfall (Table 1)	Model result (classification of current evidence); relative cost not quantified	Order of actions; programme office	Request for information, market consultation, supplier audit
Ground analogs build transferable capability only when evidence is recorded	Evidence-transfer register (Table 4)	Proposal	Facility agreements; facility operator and later flight programme	Recording requirements for configuration, loads, methods and results

The POMDP formulation provides a compatible route from the readiness profile to sequential decision-making. Its objective function can address the cost and timing of the next action while retaining capability states and mandatory acceptance conditions as distinct elements of the decision model. The main limitation is therefore not the absence of a scalar ecosystem-readiness index, but the current lack of calibrated transition and observation probabilities. These probabilities would need to be developed from programme records, supplier information, test outcomes, and comparable implementation histories.

The accompanying implementation includes checks for resource accounting, structure-first procurement precedence, quarter-level milestone handling, propagation of unknown inputs, and the separation of scenario margins from evidence-supported dates. Source and input hashes link figures and tables to the run that produced them. These measures support computational correctness and traceability; empirical validation requires comparison with configuration-specific evidence and agreed assessment criteria [33].

9 Conclusions

This study links ECLSS architectural choices to the industrial and programme capabilities required for their implementation. The HUMANS and lunar-lander cases are assessed through resource breakdowns, explicit programme-cost boundaries, TRIS-based maturation paths, and case-specific implementation roadmaps. An ecosystem-readiness profile records the available evidence, remaining work, and key dependencies, using cost and schedule as figures of merit. Essential uncertainties are retained explicitly rather than being collapsed into a single readiness score or success probability.

The scenarios show that procurement knowledge and implementation effort can affect a design trade in different ways. Current evidence supports conditional development pathways and identifies priorities for further assessment, while configuration-specific delivery conditions and cost preferences remain to be established through supplier engagement, verification planning, and programme-level decisions.

HUMANS also provides a prospective means of developing process knowledge, integration practice, operational evidence, and specialist expertise relevant to subsequent flight development. The proposed transfer

register distinguishes between knowledge that can be established through ground activities and verification that must remain flight-specific. It therefore provides a concrete basis for defining the facility's contribution to future capability, subject to the requirements of the receiving programme and the evidence ultimately obtained.

For programme and funding bodies, the profile shows, before development funds are committed, whether a gap calls primarily for information, investment, or infrastructure (Table 1), and how the resulting implementation path compares with its required milestone.

Acknowledgements

This work draws on the HUMANS analog under development at Politecnico di Torino in the IDEAS interdepartmental centre.

References

[1] L. Vega. Environmental control and life support (ECLS) systems. In L. R. Young and J. P. Sutton, editors, *Handbook of Bioastronautics*, pages 69–82. Springer International Publishing, 2021. URL https://ntrs.nasa.gov/citations/20210010644.

[2] European Space Agency. Explore2040: The european exploration strategy. Strategy report, European Space Agency, 2024. URL https://esamultimedia.esa.int/docs/HRE/Explore_2040.pdf.

[3] G. Narducci, R. Fusaro, and N. Viola. Technology roadmap methodology and tool upgrades to support strategic decision in space exploration. *Aerospace*, 12 (8):682, 2025. doi: 10.3390/aerospace12080682. URL https://doi.org/10.3390/aerospace12080682.

[4] N. Viola, R. Fusaro, V. Vercella, and G. Saccoccia. Technology roadmapping strategy, tris: Methodology and tool for technology roadmaps for hypersonic and re-entry space transportation systems. *Acta Astronautica*, 170: 609–622, 2020.

[5] M. Hekkert, R. Suurs, S. Negro, S. Kuhlmann, and R. Smits. Functions of innovation systems: A new approach for analysing technological change. *Technological Forecasting and Social Change*, 74(4):413–432, 2007. doi: 10.1016/j.techfore.2006.03.002. URL https://doi.org/10.1016/j.techfore.2006.03.002.

[6] R. Adner. Ecosystem as Structure. *Journal of Management*, 43(1):39–58, 2017. doi: 10.1177/0149206316678451. URL https://doi.org/10.1177/0149206316678451.

[7] R. Adner and R. Kapoor. Innovation ecosystems and the pace of substitution: Re-examining technology S-curves. *Strategic Management Journal*, 37(4):625–648, 2016. doi: 10.1002/smj.2363. URL https://doi.org/10.1002/smj.2363.

[8] NASA. NASA Systems Engineering Handbook. Technical Report NASA/SP-2016-6105 Rev2, National Aeronautics and Space Administration, 2017. URL https://www.nasa.gov/reference/appendix-g-technology-assessment-insertion/. Appendix G: Technology Assessment/Insertion.

[9] J. Witt, S. Hovland, D. Laurini, C. Matthias, F. Boettcher, T. Bevilacqua, and C. Redondo. On-orbit testing of the advanced closed loop system ACLS. Technical Report ICES-2020-510, European Space Agency, 2020. URL https://hdl.handle.net/2346/86479. ESA technical report presented at the International Conference on Environmental Systems.

[10] C. Tamponnet, C. J. Savage, P. Amblard, J. C. Lasserre, J. C. Personne, and J. C. Germain. Water recovery in space. Technical Report ESA Bulletin 97, European Space Agency, March 1999. URL https://www.esa.int/esapub/bulletin/bullet97/tamponne.pdf.

[11] M. S. Anderson, M. K. Ewert, and J. F. Keener. Life support baseline values and assumptions document. Technical Report NASA/TP-2015-218570/REV1, NASA, 2018. URL https://ntrs.nasa.gov/citations/20180001338.

[12] [To complete: iDREAM lander-sizing reference cited for trajectory, power and thermal modelling], 2026.

[13] J. A. Levri, J. W. Fisher, H. W. Jones, A. E. Drysdale, M. K. Ewert, A. J. Hanford, J. A. Hogan, J. A. Joshi, and D. A. Vaccari. Advanced life support equivalent system mass guidelines document. Technical Report NASA/TM-2003-212278, NASA, 2003. URL https://ntrs.nasa.gov/citations/20040021355.

[14] European Space Agency. Moon village: Conceptual design of a lunar habitat. CDF Study Report CDF-202(A), Issue 1.1, ESA Concurrent Design Facility, 2020. ACLS for four crew, 850 kg. [*Entry added in revision*].

[15] Eagle Engineering, Inc. Lunar base scenario cost estimates. Technical Report Eagle Engineering Report 88-215, NASA Lyndon B. Johnson Space Center, 1988. URL https://ntrs.nasa.gov/citations/19890005915. NASA Contract NAS9-17878; crew-module ECLSS estimate, pp. 12, 14 and 143.

[16] J. F. Connolly. After LM: NASA lunar lander concepts beyond Apollo. Technical Report NASA/SP-2020-220338, NASA Johnson Space Center, 2020. URL https://ntrs.nasa.gov/citations/20190031985. [*Entry added in revision; please confirm the 45.862 t ESAS value against the page*].

[17] C. Stromgren, C. Burke, J. Cho, W. Cirillo, A. Owens, and D. Howard. Regenerative ECLSS and logistics analysis for sustained lunar surface missions. In *IEEE Aerospace Conference*, 2022. URL https://ntrs.nasa.gov/citations/20210022453. Table 1, sub-system masses including critical spares. [*Entry added in revision; add the IEEE DOI*].

[18] A. J. Hanford. Exploration life support baseline values and assumptions document. Technical Report NASA CR-2006-213693, NASA Johnson Space Center, 2006. [*Entry added in revision*].

[19] R. Fusaro, N. Viola, D. Ferretto, V. Vercella, V. Fernandez Villace, and J. Steelant. Life cycle cost estimation for high-speed transportation systems: R. fusaro et al. *CEAS Space Journal*, 12(2):213–233, 2020.

[20] H. W. Jones. The life cycle cost (LCC) of life support recycling and resupply. Technical Report ICES-2015-295, NASA Ames Research Center, 2015. URL https://ntrs.nasa.gov/citations/20160001262.

[21] M. B. Abney, W. Schneider, B. Brown, C. Stanley, K. Lange, J. Wetzel, R. Morrow, and R. Gatens. Comparison of exploration oxygen recovery technology options using esm and lsmac. Technical Report ICES-2020-108, NASA, 2020. URL https://ntrs.nasa.gov/citations/20205004679. NASA technical paper; TRL distinction in nomenclature and Section II.B.

[22] N. P. Haggerty, B. C. White, A. Paz, D. O'Connor, N. Azim, and J. Captain. Environmental testing of a fully automated carbothermal reactor for lunar oxygen production. Technical Report ICES-2025-200, Sierra Space and NASA, 2025. URL https://ntrs.nasa.gov/citations/20250004151. 54th International Conference on Environmental Systems; TRL statement on p. 3.

[23] L. Kaldon and A. Presby. Fission surface power project. Technical Report NTRS 20240003580, NASA Glenn Research Center, 2024. URL https://ntrs.nasa.gov/api/citations/20240003580/downloads/240423%20FSP%20SPW%20Final.pdf. Space Power Workshop presentation, 23–25 April 2024; electrolysis state of the art, slide 6.

[24] NASA. Nasa systems engineering handbook: Section 6.8, decision analysis. Technical Report NASA/SP-2016-6105 Rev2, National Aeronautics and Space Administration, 2017. URL https://www.nasa.gov/reference/6-8-decision-analysis/.

[25] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X. URL https://doi.org/10.1016/S0004-3702(98)00023-X.

[26] NASA. As artemis moves forward, nasa picks spacex to land next americans on moon, April 2021. URL https://www.nasa.gov/news-release/as-artemis-moves-forward-nasa-picks-spacex-to-land-next-americans-on-moon/. Release 21-042, 16 April 2021; original award announcement.

[27] NASA. Nasa selects blue origin as second artemis lunar lander provider, May 2023. URL https://www.nasa.gov/news-release/nasa-selects-blue-origin-as-second-artemis-lunar-lander-provider/. Release 23-056, 19 May 2023; original award announcement.

[28] NASA. Report regarding nasa human landing system program. Technical report, National Aeronautics and Space Administration, April 2022. URL https://www.nasa.gov/wp-content/uploads/2023/07/human-landing-system-30-day-report.pdf. Report submitted in response to the Consolidated Appropriations Act, 2022; p. 3.

[29] NASA Office of the Chief Financial Officer. NASA FY25 Inflation Tables—to be Utilized in FY26, 2025. URL https://www.nasa.gov/ocfo/ppc-corner/ppc-models-tools/. Public draft workbook; actuals through September 2025; FY2026 projected. Inflation Table, cell BR38: 1988–2026 factor 2.99445648. Accessed 21 September 2026.

[30] S. Amalfitano, C. Levantesi, L. Garrelly, D. Giacosa, F. Bersani, and S. Rossetti. Water quality and total microbial load: A double-threshold identification procedure intended for space applications. *Frontiers in Microbiology*, 9:2903, 2018. doi: 10.3389/fmicb.2018.02903. URL https://doi.org/10.3389/fmicb.2018.02903.

[31] A. H. Ridley, C. A. Brown, J. D. Garr, II, L. L. Gavin, D. M. Hornyak, K. P. Toon, P. A. Caradec, and A. Williams. International Space Station as a Testbed for Exploration Environmental Control and Life Support Systems - 2023 Status. Technical Report ICES-2023-259, National Aeronautics and Space Administration, 2023. URL https://ntrs.nasa.gov/citations/20230006553.

[32] A. Saltelli. Composite indicators between analysis and advocacy. *Social indicators research*, 81(1):65–77, 2007.

[33] NASA. Nasa handbook for models and simulations: An implementation guide for nasa-std-7009. Technical Report NASA-HDBK-7009A, National Aeronautics and Space Administration, 2019. URL https://standards.nasa.gov/sites/default/files/standards/NASA/A/0/nasa-hdbk-7009a.pdf.