POLITECNICO DI TORINO
Repository ISTITUZIONALE

From Strategic Technology Gaps to Architectural Trade-Offs: An Integrated Systems Engineering Framework

Original
From Strategic Technology Gaps to Architectural Trade-Offs: An Integrated Systems Engineering Framework / Festa, L.M., Corpino, S., Stesina, F.. - ELETTRONICO. - (2026). (77rd International Astronautical Congress (IAC) Antalya, Türkiye 5-9 October, 2026).

Availability:
This version is available at: 11583/3017161 since: 2026-10-28T16:18:50Z

Publisher:
IAF Astro

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
IAC/IAF postprint versione editoriale/Version of Record

Manuscript presented at the 77rd International Astronautical Congress (IAC), Antalya, Türkiye, 2026. Copyright by IAF

(Article begins on next page)

27 November 2026

77th International Astronautical Congress (IAC 2026), Antalya, Türkiye, 5–9 Oct 2026.
Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

IAC-26,D1,IP,20,x113892

From Strategic Technology Gaps to Architectural Trade-Offs: An Integrated Systems Engineering Framework

Mr. Leonardo Maria Festa[a*], Prof. Sabrina Corpino[a], Dr. Fabrizio Stesina[a]

[a] *Department of Mechanical and Aerospace Engineering, Politecnico di Torino, Italy, name.surname@polito.it*

** Corresponding author, leonardomaria.festa@polito.it*

Abstract

Early-phase systems engineering for the New Space economy must connect commercial and institutional service needs with technical design decisions. This paper presents an iterative four-phase framework linking strategic technology gaps, mission opportunities, Operational Capabilities (OCs), architecture trade-offs, and technology-application priorities. First, capability gaps from the Global Exploration Roadmap of the International Space Exploration Coordination Group are weighted and related to a mission portfolio. Second, mission decomposition identifies OCs and an evaluation framework comprising Key Performance Parameters, constraint margins, and strategic drivers. Mission relevance and distinct mission–OC coverage determine normalized OC priorities. Third, morphological analysis generates and screens candidate architectures before parametric baselines are compared using weights from the Analytic Hierarchy Process and Multi-Criteria Decision Analysis. Fourth, the selected architecture is decomposed into technology applications whose relative contributions are weighted by the OC priorities. A lunar Space Rover Platform case demonstrates the process: 24 concepts are generated, 15 pass context screening, and six representative architectures are baselined. The small, distributed, modular architecture ranks first under the nominal assumptions. Distributed-system management receives its highest relative technology contribution score, followed by onboard software and avionics. Sensitivity analysis distinguishes stability under category-weight changes from dependence on engineering judgments. With the nominal stored-energy penalty retained, removing the assumed fleet thermal-resilience credit narrowly favors the large modular alternative; removing both adjustments leaves the distributed modular concept first. This interaction identifies coupled assumptions requiring subsequent energy, thermal, and operational assessment. The framework makes mission needs, scoring judgments, and their consequences explicit, providing a practical basis for reviewing early architecture choices and directing further engineering work.

Keywords: *systems engineering; space systems; lunar exploration; methodology; multi-criteria decision analysis*

Acronyms and abbreviations

AHP	Analytic Hierarchy Process
CI	Consistency Index
CR	Consistency Ratio
DoD	Depth of discharge
GER	Global Exploration Roadmap
KPP	Key Performance Parameter
MCDA	Multi-Criteria Decision Analysis
OC	Operational Capability
SA	System architecture
SRP	Space Rover Platform
TA	Technology Application

1. Introduction

The New Space economy motivates early system definition around recurring services, cost-conscious development, operational autonomy, maintainability, and growth across missions. Institutional programmes and commercial providers participate in this evolving context [1, 2]. For the framework developed here, these drivers are part of the decision problem: they inform the mission portfolio, morphological alternatives, screening priorities, and strategic evaluation criteria. Institutional demand can support commercially provided services: ESA's Moonlight initiative, for example, envisages the agency as an anchor customer while providers can serve additional customers [3].

A lunar Space Rover Platform (SRP) provides a useful application because payload hosting and surface services can be realized by different combinations of scale, composition, modularity, and autonomy. Modular rover demonstrators provide practical examples of configurable subsystem and payload arrangements [4]. NASA's Cooperative Autonomous Distributed Robotic Exploration (CADRE) project develops a team of three small rovers for autonomous cooperative exploration of the lunar surface [5]. Its development motivates the study of distributed rover configurations and the early comparison of their coordination, service, and resource demands. Institutional needs also remain relevant: the Global Exploration Roadmap (GER) of the International Space Exploration Coordination Group (ISECG) identifies exploration capabilities and technology gaps [6]. Translating these inputs into a system concept requires judgments about mission relevance, operational capabilities, feasible alternatives, and the trade between physical performance and programmatic value.

Structured support for these decisions has substantial precedent. The NASA Systems Engineering Handbook specifies alternatives, criteria, evaluation, and the

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

recording of assumptions and uncertainty [7]. Lafleur et al. combine AHP objective prioritization, Quality Function Deployment (QFD)-like instrument-to-objective mappings, and architecture downselection [8]. Multi-Attribute Tradespace Exploration (MATE) connects stakeholder preferences with model-based architecture evaluation, sensitivity, and higher-fidelity design; it also explicitly addresses learning about system drivers [9]. These precedents establish the methodological basis on which the present workflow builds.

Other approaches resolve specific parts of the decision problem. Mosleh et al. evaluate monolithic and distributed architecture value under uncertainty [10], while Chen et al. couple commercialization decisions with space-logistics architecture [11]. These analyses provide useful reference points but answer different questions from the present workflow. Here, commercial drivers define the evaluation context and qualitative criteria; stochastic reliability and economic optimization remain outside the case model.

This paper connects early-definition decisions for a New Space service context: strategic technology gaps are related to mission opportunities; mission activities establish operational capabilities and evaluation criteria; morphological candidates are screened and compared through parametric baselines; and the capabilities are reused to prioritize technology applications for the selected concept. This organization makes the interfaces between decisions explicit and retains information useful when an assumption is revisited. Established AHP and MCDA methods supply the numerical tools [12, 13]; the contribution is their documented application across these interfaces. In particular, the entry point is a technology-gap-based mission portfolio, and the output extends beyond architecture preference to capability-linked technology priorities. Retaining these links allows a design team to trace a preferred concept back to mission needs and identify which assumptions merit further engineering analysis.

The framework supports comparison and *capability-evolution scenario analysis*: conditional reassessment of which concepts merit attention when an assumed capability or cost assessment changes. Preference sensitivity asks whether a choice survives different priorities; technology scenarios ask which assumed improvements would change that choice. These uses can identify decision-driving judgments and candidate development targets for subsequent physical and economic assessment.

The paper illustrates this approach through six lunar SRP baselines spanning standalone and distributed configurations, integrated and modular designs, and two autonomy levels. The first sections explain the methodology and its application, including technology priorities for the nominally preferred distributed modular concept. A later section demonstrates bounded preference and judgment sensitivity and discusses candidate-set dependence. The evaluation focuses on traceability and reassessment within a conceptual study; its modeling assumptions and validation needs are discussed in Section 5.3.

2. Integrated methodology

2.1 System boundary and decision context

A system architecture can be compared meaningfully only after the System of Interest (SoI) and its boundary have been stated. The boundary determines which behaviors must be supplied by the system and which external entities instead impose interfaces, constraints, or enabling services. The proposed process begins by: (1) naming the SoI and its top-level purpose; (2) identifying interfacing systems and operational environments across the lifecycle; and (3) classifying every interaction across the boundary. This distinction later separates performance criteria associated with system functions from constraint margins associated with the external context.

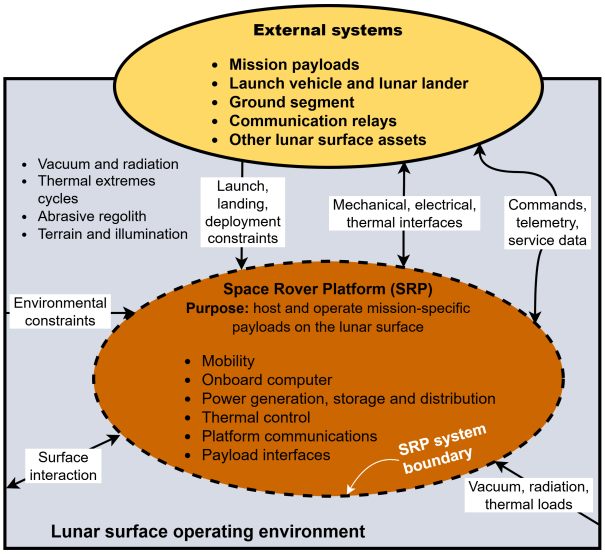

Fig. 1. SRP system-boundary view distinguishing platform functions, external interfaces, and environmental constraints.

For the case study, the SoI is the SRP, shown in Fig. 1. Its top-level purpose is to host and operate mission-specific payloads on the lunar surface. Mobility, onboard processing, power, thermal control, communications, and the payload interfaces needed for basic platform operation are inside the boundary. Mission payloads remain external so that different services can be compared without presupposing a single instrument suite. The launch vehicle, lunar lander, ground segment, communication relays, other surface assets, and the lunar environment, including abrasive regolith [14], are also external. They create mechanical, electrical, thermal, communication, and operational interfaces that become constraints or non-functional requirements.

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

2.2 Four-phase traceability chain

The framework is organized as four transformations. Phase 1 establishes the strategic foundation by selecting relevant roadmap gaps, prioritizing them, and propagating the priorities to a portfolio of mission statements. Phase 2 decomposes the mission lifecycle to identify recurring OCs and derives criteria able to discriminate between architectures. Phase 3 generates an explicit morphological design space, removes invalid combinations, creates comparable parametric baselines, and executes the architectural trade-off. Phase 4 decomposes the selected architecture into TAs and prioritizes them against the OCs established earlier. These artifacts support subsequent phases through both calculations and engineering derivations. Mission decomposition informs criterion selection, for example, but does not automatically determine architecture-category weights.

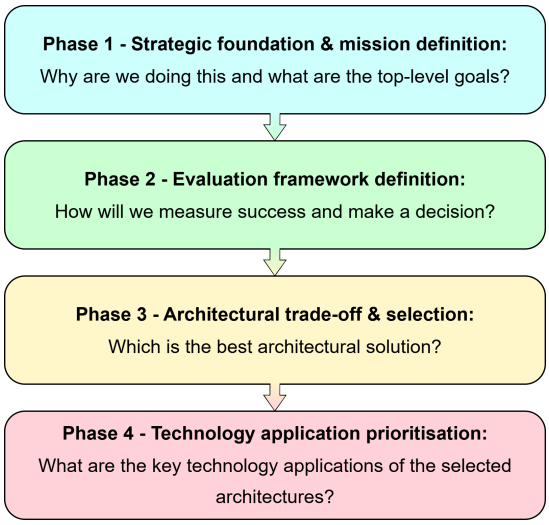

Fig. 2. Four-phase framework and its guiding questions. "Best" denotes the preferred evaluated concept under stated assumptions. Arrows indicate information flow; reassessment can return to any phase.

2.3 Mathematical backbone

Expert judgment is necessary in early design, but it should be visible, internally consistent, and revisable. AHP is used only to derive weights, avoiding its direct use as the final architecture selector. For a reciprocal pairwise comparison matrix of order n, logical consistency is assessed through the Consistency Index (CI) and Consistency Ratio (CR):

$$\mathrm{CI} = \frac{\lambda_{\max} - n}{n-1}, \qquad \mathrm{CR} = \frac{\mathrm{CI}}{\mathrm{RI}}, \tag{1}$$

where $\lambda_{\max}$ is the maximum eigenvalue and RI is Saaty's Random Index. Matrices are accepted when $\mathrm{CR} < 0.10$ [12]. The phase diagrams summarize the procedure. Their "cross-impact matrix" labels denote the relationship matrices in Phases 1 and 4 and the architecture–criterion evaluation matrix in Phase 3.

The use of scored relationships to connect decision layers follows the mapping logic of Quality Function Deployment (QFD) [15]. The present implementation uses relationship matrices to propagate priorities within the quantified stages.

In Phase 1, if w_i is the normalized priority of source element i, and r_{ij} is the scored relationship between source element i and target element j, the target priority used in this work is

$$p_j = 100\frac{\sum_i w_i r_{ij}}{\max_k \left(\sum_i w_i r_{ik}\right)}. \tag{2}$$

The scoring scale must be defined before assigning values. Phase 2 converts the mission decomposition into a binary incidence I_{mk}, equal to one when mission m requires OC k, irrespective of repeated activity rows. If p_m is the Phase 1 mission score, the normalized mission relevance α_m, strategic coverage q_k, and OC priority v_k are

$$\alpha_m = \frac{p_m}{\sum_{m'} p_{m'}}, \quad q_k = \sum_m \alpha_m I_{mk}, \quad v_k = \frac{q_k}{\sum_{k'} q_{k'}}. \tag{3}$$

Each mission contributes its full relevance to every OC it requires: the resulting weights represent strategic coverage rather than probability, workload, or allocation of a finite mission budget. Because a mission can require several OCs, $\sum_k q_k$ can exceed one; $\sum_k v_k = 1$. Phase 4 normalizes each TA–OC relationship column before aggregating the adopted OC priorities:

$$T_t = \sum_k \left(\frac{r_{tk}}{\sum_{t'} r_{t'k}}\right) v_k, \qquad S_t = 100\frac{T_t}{\max_{t'} T_{t'}}. \tag{4}$$

Here, r_{tk} is the judgment-based relationship score, T_t is a weighted relative contribution, and S_t is its maximum-scaled score.

The final MCDA separates criteria into KPP, constraint, and strategic categories. Let $\hat{s}_{aj}$ be the normalized score of architecture a against criterion j, and let W_c be the AHP-derived weight of category c. With equal criterion importance inside each category, the global utility is

$$U(a) = \sum_c W_c \left(\frac{1}{N_c}\sum_{j\in c} \hat{s}_{aj}\right), \tag{5}$$

where N_c is the number of criteria in the category. Beneficial raw values are divided by their criterion-column sum; recharge time is first transformed to its reciprocal so that shorter times receive higher scores. AHP weights are obtained by column normalization and row averaging. Some baseline and criterion calculations are linked, while

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

Table 1. Traceability artifacts generated by the four-phase framework.

Phase	Input	Transformation	Output	SRP instantiation
1	Strategic roadmap and SoI boundary	AHP prioritization and gap-mission relationship matrix	Weighted mission portfolio	16 GER gaps related to 11 lunar missions
2	Missions, Phase 1 scores, and lifecycle context	Binary mission–OC incidence, mission-weighted OC priorities, and criteria definition	OCs and weighted evaluation framework	12 OCs; 44 distinct mission–OC pairs; 5 KPPs, 5 constraint criteria, 6 strategic drivers
3	Morphological options and evaluation framework	Context screening, representative baselines, and MCDA	Conditional architecture preference	24 generated, 15 retained, and 6 baselined architectures
4	Selected architecture, TA–OC relationships, and OC priorities	Architecture decomposition, column normalization, and weighted aggregation	Relative TA contribution scores	11 TAs demonstrated for the distributed modular concept

other interfaces contain supplied judgments or copied inputs. Scenario analysis therefore requires identifying and updating every affected relationship. Equal weights within each category provide the reference preference structure for this application, assigning each criterion the coefficient W_c/N_c. Separate criterion-level preferences were not elicited.

2.4 Decision insight and scenario use

The output of the framework is more than a winning identifier. Strategic assessment reveals why a mission is relevant; mission decomposition exposes required capabilities and candidate evaluation measures; architecture exploration identifies contenders, exclusions, and assumptions that drive their ranking; and TA prioritization locates technology contributions across the capability set. These artifacts support review of the rationale and targeted reassessment when assumptions change.

Three kinds of update must be distinguished. Preference sensitivity changes weights while retaining candidate attributes. Technology-evolution scenarios change specified capability, feasibility, or cost assumptions and recalculate the affected scores. Context changes may require revising missions, constraints, or even the candidate set. Section 6 discusses reassessment after the main application. A realistic development scenario must also revise coupled performance, mass, cost, and risk assumptions where the proposed technology affects them.

3. From strategic gaps to evaluation criteria

3.1 Strategic gaps and mission portfolio

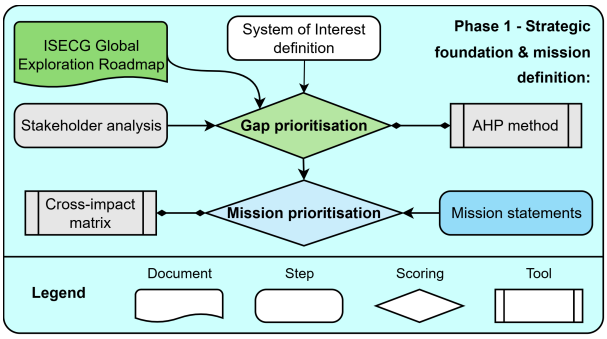

Fig. 3. Phase 1: strategic-gap prioritization and mission-portfolio assessment. The relationship matrix makes the basis of mission relevance explicit.

The GER was screened against the SRP boundary to identify gaps for which a surface robot is either a direct technology target, an enabling user, or a useful demonstrator. Sixteen gaps were retained across landing, autonomous systems, navigation, power, transversal technologies, and surface robotics. A pairwise AHP comparison then ranked their importance for the research and commercial context of the case study. The resulting matrix satisfied the consistency requirement with $\mathrm{CR} = 0.0648$. Surface mobility and exploration was used as the normalized reference at 100.0, followed by autonomous vehicle system management at 94.6 and dust mitigation at 88.9. These priorities reflect the context of the decision maker; they are structured judgments rather than universal technology rankings.

The strategic inputs were next related to 11 mission statements covering exploration, resource and construction services, operations support, logistics and crew

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

Table 2. Leading Phase 1 outputs for the SRP case.

Ranked element	Score
Technology gaps	
Surface mobility and exploration	100.00
Autonomous vehicle management	94.60
Dust mitigation	88.90
Mission statements	
Technology validation	100.00
Terrain reconnaissance and mapping	72.19
Site and infrastructure development	61.49
Regolith and volatile utilization	61.13

support, and cross-cutting technology services. Each gap-mission relationship was scored on a predefined 0–9 scale, from no relationship to an enabling or indispensable relationship. Equation (2) combined this matrix with the gap weights. Technology validation ranked first at 100.0, because it can both exercise and mature many of the selected gaps. Terrain reconnaissance and mapping ranked second at 72.19, followed by site and infrastructure development at 61.49 and regolith and volatile utilization at 61.13. The ranking indicates strategic relevance, not temporal precedence: a lower-ranked reconnaissance activity may still be required before a higher-ranked resource mission can be attempted.

3.2 Mission lifecycle and operational capabilities

A mission statement is too broad to serve directly as an architecture criterion. Phase 2 therefore decomposes the generic lifecycle into pre-launch, launch, cruise, landing and deployment, commissioning, lunar-surface operations, end of mission, and end of life. The common phases expose constraints imposed by the launcher, lander, deep-space transfer, and lunar environment. The surface-operations phase is then decomposed separately for each mission statement to identify the platform behavior required by each operational step.

This analysis produced 12 unique OCs from 46 activity-level occurrences. After deduplicating repeated activity rows, the result is 44 distinct mission–OC pairs. Traversal/navigation and relay communications each span all 11 missions and have a mission-weighted priority of 0.2498; sensing and data collection spans eight missions and has a priority of 0.1889. Technology validation appears in only one mission but has a priority of 0.0369 because that mission has the highest Phase 1 relevance. The priority therefore combines strategic relevance with breadth; it is neither an operational frequency nor a workload estimate. The same OC list is reused in Phase 4 to weight the column-normalized TA relationships and as a functional-completeness check for the selected architecture.

3.3 Evaluation framework

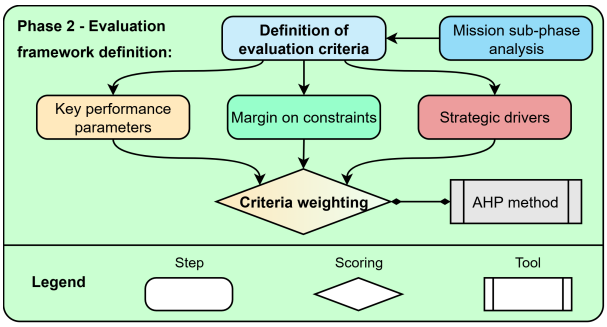

Fig. 4. Phase 2: mission analysis informs the evaluation criteria. AHP weights the three categories; criteria within each category are equally weighted.

Criteria are selected for their expected ability to distinguish architectures. For example, all viable rovers must communicate, but at this stage the communication implementation is not defined precisely enough to produce a useful differentiating KPP. Architecture scale and composition, by contrast, create measurable differences in total payload capacity, energy storage, traversable range, area coverage, and margins against external constraints.

The resulting evaluation framework contains three complementary categories. KPPs capture quantitative platform performance. Margin scores distinguish candidates after applicable mandatory requirements are screened; a weighted score cannot compensate for a mandatory failure. The present early study provides context screening rather than complete technical qualification. Strategic drivers represent programmatic attributes required by the intended service context. Table 3 records the complete set used by the case study.

The category weights describe the priorities adopted for this research-stage service-platform study: constraint assessments receive the greatest weight, followed by quantitative performance and strategic drivers. The strategic category's approximately 20% weight is a declared preference, not a measurement of commercial relevance. Service breadth, modularity, autonomy, and programmatic suitability also shape the mission portfolio, morphological choices, and screening. Other decision contexts can use different category preferences within the same evaluation structure.

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

Table 3. Architecture evaluation framework derived in Phase 2.

Category	Weight	Criteria
Key Performance Parameters	31%	Maximum coverable area; maximum reachable distance; total surface payload capacity; usable battery energy per unit; recharge time (shorter preferred).
Constraint margins	49%	Lander-capacity utilization; stored-power margin; operational power margin; thermal resilience during a 100-h dark period; structural robustness against a major mechanical failure.
Strategic drivers	20%	Lifecycle cost; reliability and redundancy; scalability; longevity and maintainability; development risk; operational simplicity.

4. Architecture generation and trade-off

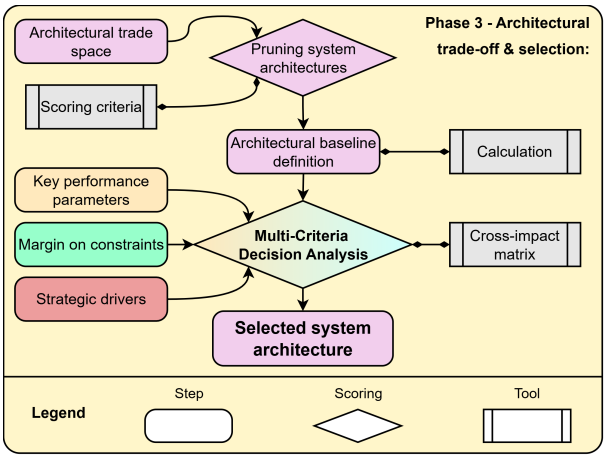

Fig. 5. Phase 3: candidate screening, baseline definition, and architecture comparison. The matrix at the MCDA step is the architecture–criterion evaluation matrix; screening uses separate viability scores.

4.1 Morphological design space and screening

Morphological analysis structures candidate generation through combinations of explicitly defined design dimensions [16]. Here, scale, composition, integration philosophy, and autonomy define the discrete alternatives in Table 4, giving $2\times2\times2\times3 = 24$ theoretical combinations.

Table 4. Morphological dimensions used to generate candidate architectures.

Dimension	Options
Scale	Large (at least 500 kg); small (up to 150 kg)
Composition	Standalone; distributed (at least three units)
Philosophy	Integrated; modular payloads and/or subsystems
Autonomy	Level 0 teleoperation; Level 1 supervised functions; Level 2 mission supervision

Context-specific screening rules removed nine combinations before numerical scoring. Large distributed systems were excluded because their launch and operational burden was inconsistent with the near-term commercial context. Distributed Level 0 concepts were excluded because direct teleoperation of multiple units was not operationally credible. A large standalone modular Level 0 concept was also removed because it added interface complexity without autonomy able to exploit reconfiguration.

The remaining 15 candidates were screened for technical feasibility, programmatic viability, and operational suitability, weighted at 20%, 35%, and 45%, respectively. Each morphological option carries a score in each category. For architecture a, the screening score is

$$V_a = \sum_{c\in\{T,P,O\}} w_c \sum_{d} s_{d,o(a,d),c}, \tag{6}$$

where $o(a,d)$ is the option selected in dimension d. This additive model makes option-level judgments explicit but omits interactions between options. Scores are viability assessments, not measured probabilities or costs.

The four highest screening scores are 10.00 for SA-SDM-2, 9.95 for SA-SSM-2, 9.85 for SA-SDI-2, and 9.80 for SA-SSI-2. The comparison includes these four concepts, SA-LSM-2 to represent a large modular alternative, and SA-LSI-1 as the large integrated reference. Thus six architectures are baselined: five alternatives plus one reference. This selection retains both leading screened concepts and contrasting large-rover designs.

After the SA prefix, the identifiers encode scale (S: small; L: large), composition (S: standalone; D: distributed), and philosophy (M: modular; I: integrated). The final number denotes autonomy level. Thus SA-SDM-2 denotes a small, distributed, modular architecture with Level 2 mission supervision. The shortened labels SDM, SSM, SSI, LSI, LSM, and SDI retain these three architectural attributes. The two distributed concepts therefore allow the integration philosophy to be examined within a fleet configuration, alongside the standalone counterparts.

4.2 Parametric architecture baselines

Abstract labels such as “large” or “distributed” are insufficient for scoring. Each shortlisted concept was therefore assigned a common set of preliminary mass, power, thermal, and mobility assumptions. These values are intended to be comparable Phase 0/A estimates, not subsystem specifications.

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

Table 5. Decision-driving baseline values for six evaluated architectures (five alternatives and a reference).

Architecture	Surface mass (kg)	Lander use (%)	Payload (kg)	Unit energy (Wh)	Recharge (h)	Range (km)	Area (km²)
SA-SDM-2	840	56.0	75.6	8249	40.41	536.54	7.51
SA-SSM-2	144	9.6	13.0	9764	47.83	529.24	1.06
SA-SSI-2	121	8.1	13.2	8923	43.71	527.59	1.06
SA-LSI-1	550	36.7	75.0	24411	79.72	317.55	1.59
SA-LSM-2	840	56.0	105.0	34007	41.64	315.98	1.58
SA-SDI-2	840	56.0	69.3	8923	43.71	527.59	7.39

SA-SDM-2 is represented by seven 100-kg units with a 20% modular-system mass allowance: $7\times100\times1.20 = 840$ kg. The seven-unit fleet was selected to match the SA-LSM-2 surface mass, which was defined from representative large-rover masses while retaining margin against the Argonaut-class lander capacity, thereby preserving an approximately mass-equivalent comparison. SA-SDI-2 uses seven 110-kg integrated-unit analogues with battery, structure, and avionics fractions of 0.50, 0.30, and 0.11. The remaining 0.09 gives 9.9 kg payload per unit and 69.3 kg across the fleet. A separate 9.1% system mass allowance gives $7\times110\times1.091 = 840.07$ kg, rounded to 840 kg in Table 5. The allowance is a preliminary allocation without a subsystem breakdown; solar and mobility inputs use the same assumptions as SA-SSI-2. Coordination hardware and power are not separately sized. Mass and payload values are preliminary allocations for comparative evaluation, with assumptions retained in the supporting calculations and subsystem budget closure reserved for detailed design. The lander reference capacity was 1500 kg, consistent with the Argonaut-class context [17]. Battery capacity was estimated from a common lithium-ion (Li-ion) 21700 cell model with 70% usable depth of discharge (DoD). Solar generation assumed lunar irradiance of 1361 W/m², 30% conversion efficiency, and a 0.5 orientation factor. Thermal assessment uses an available-power budget of $E_{\mathrm{av}}/100$ h for a 100-h dark interval. It does not calculate heat losses or prove survival. Actual south-polar shadow intervals are site-dependent; for example, the Volatiles Investigating Polar Exploration Rover (VIPER) mission planning defined safe havens using at most 50 h of darkness while Earth was below the horizon [18]. Mobility calculations assumed a common terrain model with a 30-degree limiting slope over 20% of the traverse and a rolling-resistance coefficient of 0.2.

For architecture a, usable battery energy and maximum traverse range were computed as

$$E_{\mathrm{av},a} = \mathrm{DoD}\, E_{\mathrm{tot},a}, \qquad R_a = \frac{E_{\mathrm{av},a}}{e_{\mathrm{km},a}}, \tag{7}$$

where $e_{\mathrm{km},a}$ is the calculated electrical energy consumed per kilometre. For a distributed architecture, nominal mapped area was estimated by

$$A_a = n_a R_a b_a, \tag{8}$$

where n_a is the number of units and b_a is the effective lateral coverage width. This system-level parallelism is the primary performance mechanism by which a distributed architecture differs from a single rover.

Table 5 retains the baseline quantities that most directly explain the trade. "Lander use" is the fraction of the 1500-kg capacity occupied by the deployed system. Energy, recharge, and range refer to an individual active unit; mass and payload refer to the surface system, and area aggregates the distributed units. Consequently energy is a per-unit reserve criterion rather than total fleet storage. LSM provides the largest total payload and per-unit energy reserve; SDM carries approximately the same total payload as LSI and has greater nominal coverage. Retention of useful capability after a unit loss is assumed in the qualitative robustness assessment, not independently simulated.

4.3 Six-architecture MCDA comparison

Criterion scores use the normalization in Section 2, with recharge treated as a cost criterion. The payload criterion represents total surface payload, rather than payload fraction; its preliminary 0–10 scaling cancels under column-sum normalization. Operational power margins of 0.526–0.729 give a score of three for all six candidates under the implemented ceiling rule.

SA-SDI-2 is assigned distributed-system attributes of redundancy, scalability, and tolerance of a single-unit loss, while its integrated construction receives a maintainability score of one. Its scores for lifecycle cost and development risk are one because autonomous coordination remains an assumed enabling technology. The nominal assessment assigns SDM and SDI a thermal-resilience score of two, including a one-level allowance conditional on developing an effective fleet survival strategy. The allowance represents assumed retention of useful fleet service through the modeled dark interval. Its feasibility has not been established by a heat-loss, survival-power, or common-cause failure model. Section 6 therefore

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

evaluates the same architectures with this allowance removed. Structural robustness is scored three for SDM and two for SDI. These are explicit rubric judgments, not measured reliability, cost, or TRL. The nominal rubric also subtracts one point from the stored-energy score of both distributed concepts. This adjustment has no independently established physical basis. Section 6 therefore tests its removal both alone and together with the thermal credit to expose its influence on the result.

Table 6. Weighted category contributions for the six-architecture comparison. Final scores are scaled to the highest aggregate value.

Architecture	KPP	Margin	Strategic	Final
SA-SDM-2	0.0663	0.0944	0.0362	100.00
SA-LSM-2	0.0649	0.0862	0.0337	93.86
SA-SDI-2	0.0642	0.0855	0.0311	91.79
SA-LSI-1	0.0477	0.0773	0.0300	78.72
SA-SSM-2	0.0342	0.0780	0.0370	75.79
SA-SSI-2	0.0346	0.0691	0.0297	67.77

SA-SDM-2 ranks first and leads the aggregated KPP and margin categories; SA-SSM-2 leads the strategic category. SDM's coverage and margin assessments outweigh LSM's advantages in payload and per-unit energy under the adopted weights. The final values are scores relative to the best evaluated concept, not percentages of physical performance or probabilities of superiority. Section 6 examines the sensitivity of this conditional preference.

5. Technology priorities and methodological discussion

5.1 From architecture to technology applications

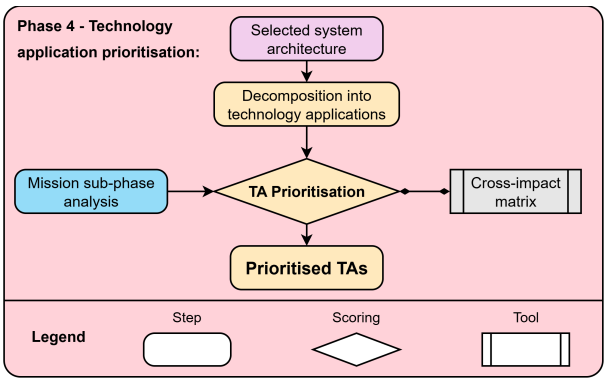

Fig. 6. Phase 4: architecture decomposition and technology-application prioritization using the mission-derived OCs. The application below concerns SA-SDM-2.

Architecture selection is not the end of the proposed process. Phase 4 asks which technical capabilities must be developed, integrated, or procured to realize the selected concept. For the nominally preferred SA-SDM-2 architecture, TA–OC relationships describe functional contribution. Mass allocations and thermal-survival assumptions enter the architecture evaluation separately. SA-SDM-2 was decomposed into 11 TAs: mobility, structure, thermal control, avionics, power, communications, sensors, onboard software, manipulators, payload interfaces, and distributed-system management.

Each TA was compared against the 12 OCs using a predefined 0–9 relationship scale. The relationships are column-normalized and weighted by the mission-relevance-weighted priorities of distinct mission–OC pairs. The result combines strength of contribution with strategic breadth across the modeled mission portfolio. In-house development, procurement, and partnership decisions additionally require additional evidence on suppliers, resources, intellectual property, and development risk.

One example connects all stages of the mission-to-technology chain. The surface mobility and exploration gap has a relationship score of nine with the terrain reconnaissance and mapping mission, whose aggregate Phase 1 score is 72.19. The mission activity "Perform Systematic Traverse" requires the traversal/navigation OC. This capability occurs in all 11 missions; aggregating their relevance gives an OC priority of approximately 0.2498. In Phase 4, mobility receives a relationship score of nine against this OC, with a relationship-column sum of 70. Its contribution to mobility's weighted TA total is therefore $(9/70) \times 0.2498 \approx 0.0321$. The example links a strategic gap to a named mission, an operational activity, a shared capability, and a technology contribution; the OC weight includes the complete mission portfolio.

Distributed-system management ranks first at 100.00, followed by onboard software at 96.72 and avionics at 95.35. These results characterize the illustrated distributed architecture: coordination, task allocation, collective navigation, fault management, and mission-level autonomy distinguish a distributed platform from several independent rovers. Communications and thermal control are closely spaced at 72.45 and 72.05, respectively, and their adjacent order should not be over-interpreted.

5.2 Methodological contribution and decision insight

The main strength of the framework is traceability. The selected TAs can be related back to OCs, mission statements, and strategic gaps through explicit artifacts and matrices. This makes the rationale inspectable and allows a team to locate the judgment responsible for a disputed outcome. The framework also treats physical performance, constraint margin, and strategic value within one decision structure. The sensitivity analysis in Section 6 demonstrates the value of retaining these interfaces by identifying how thermal and stored-energy scoring assumptions affect architecture preference.

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

Table 7. Relative technology-application contributions for nominally preferred SA-SDM-2, using mission-weighted OC priorities.

Rank	Technology application	Score
1	Distributed-system management	100.00
2	Onboard software	96.72
3	Avionics	95.35
4	Power	78.04
5	Mobility	73.88
6	Communications	72.45
7	Thermal control	72.05
8	Sensors	61.34
9	Structure	50.96
10	Payload interfaces	46.30
11	Manipulators	27.42

5.3 Assumptions and limitations

Gap weights, relationship scores, qualitative strategic assessments, and some baseline assumptions were established in a single academic case study without a broad expert elicitation campaign. The consistency ratios assess the logical coherence of pairwise comparisons; they do not independently validate the preferences. Industrial use should collect multiple expert perspectives and retain the rationale and disagreement behind score assignments.

Equal within-category criterion coefficients have not been separately elicited, and candidate-relative normalization can change aggregate scores when alternatives are added or removed. The ordinal treatment, possible overlap between criteria, and judgment-based TA relationships therefore limit interpretation. Category weighting, feasibility screening, and architecture parameter updates require explicit engineering judgments alongside the calculated mission, OC, and TA priorities.

Model fidelity is a further limitation. The baseline is deliberately simple enough to evaluate multiple concepts early, but it cannot replace detailed power, thermal, mobility, reliability, or cost models. The thermal allowance is conditional, and range and area coverage are comparative indicators under common assumptions. Relative TA contribution is not a schedule, procurement decision, investment return, or measure of development criticality.

5.4 Development guidance and transferability

Scenario reassessment provides a prospective use of the retained decision artifacts: it can identify assumptions under which a shortlist or architecture preference changes. Combined with TA–OC relationships, this can help frame a development investigation. Broad functional contribution and development criticality answer different questions. A technology with a lower TA contribution score can remain essential if its feasibility controls the architecture preference. The thermal and stored-energy sensitivity checks provide such questions and should be considered alongside the TA ranking. Development effort, achievable performance, dependencies, and opportunity cost must be considered before allocating resources. The present output is an illustrative set of technology priorities and a basis for conditional reassessment, rather than a time-phased development roadmap.

New Space drivers shape this workflow at several levels. Mission breadth motivates a platform service perspective; modularity and autonomy enter the morphological choices; programmatic and operational considerations shape the illustrative screening weights; and lifecycle cost, scalability, maintainability, development risk, and operational simplicity enter the final criteria. These choices define the service context assessed in this study; a business case would require separate economic analysis.

The reusable element is the organization of decisions, artifacts, and reassessment. Applying it to Earth observation would require new strategic inputs, mission activities, capabilities, architectural dimensions, models, and scoring judgments. The process structure can be retained and tailored, with domain-specific content and validation. Comparative studies with design teams are also needed to assess transferability, designer understanding, decision quality, or engineering effort.

6. Sensitivity and decision insight

6.1 Bounded demonstration

The retained evaluation matrices support reassessment after the main design walkthrough. Two small experiments illustrate this use: a change in category preferences and a change in the qualitative assessment of distributed thermal resilience. All six alternatives are evaluated, with normalization recomputed when raw criterion scores change. An additional check tests the stored-energy penalty and its interaction with the thermal credit. All checks retain the same physical baseline inputs.

For preference sensitivity, one category weight t varies over $[0, 1]$. The other two weights retain their nominal ratio:

$$W_k(t) = t, \qquad W_j(t) = (1-t)\frac{W_j^0}{1-W_k^0}, \quad j \neq k. \tag{9}$$

Candidate attributes and equal within-category weights remain fixed. SDM stays first throughout the KPP and margin paths. Along the strategic-weight path in Fig. 7(a), SDM remains first below 93.73%; SSM leads above this threshold; both tie at the threshold. The nominal strategic weight is 19.76%. This is stability to the specified category-weight changes, not robustness to every criterion judgment or physical uncertainty.

The thermal experiment identifies an engineering condition relevant to the architecture decision: what survival strategy and retained mission capability would justify the assumed fleet credit? With the nominal stored-energy penalty retained, removing the thermal credit narrowly favors LSM, whereas the nominal conditional

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

assessment favors SDM. This comparison identifies a development question for subsequent thermal, power, and operational analysis. The ordinal changes are judgment scenarios with fixed physical inputs; rating one does not mean all rovers fail, and rating three is not an achieved technology state.

At a hypothetical rating of three, SDM remains first and SDI moves ahead of LSM, at 91.99 versus 89.58. The lower LSM utility in Fig. 7(b) follows from the changed normalization denominator, although its own raw scores remain unchanged.

Removing the one-point stored-energy penalty raises the corresponding scores of both distributed concepts from three to four. At nominal category weights and with the thermal credit retained, SDM remains first; SDI rises to second at 91.94, ahead of LSM at 91.05. Removing both the stored-energy penalty and the thermal credit also leaves SDM first, followed by LSM at 97.39 and SDI at 91.65. Column normalization is recomputed across all six concepts in each check. Thus the winner change caused by removing the thermal credit depends on retaining the stored-energy penalty; the two rubric assumptions must be assessed together.

6.2 Interpretation and scope

Together, the experiments distinguish a stable preference under the tested weight changes from dependence on a consequential engineering judgment. The combined checks identify a question for further investigation: how should fleet energy reserves and retained mission capability determine the stored-energy and thermal-resilience assessments? The scenario scores express changes in judgment. Common-cause failures, required survival power, and coupled mass, cost, and operational penalties require separate models. A future capability-evolution study could supply those models and update the corresponding criteria, enabling technology assumptions to be compared on a physical basis.

Column-sum normalization makes criterion scores depend on the evaluated candidate set. Adding or removing alternatives can therefore change relative utilities and may reverse rankings even when the remaining alternatives retain their raw attributes. In the present check, removing SDI leaves SDM first and LSM second at 89.95, compared with 93.86 in the six-concept set. This particular outcome does not establish general invariance to candidate-set changes, and the score differences are not absolute engineering benefits.

Reassessment can enter earlier phases as well: revised strategic inputs can change mission relevance, mission decomposition can expose a missing capability, and screening judgments can reopen an excluded alternative. The present numerical example concerns final architec-

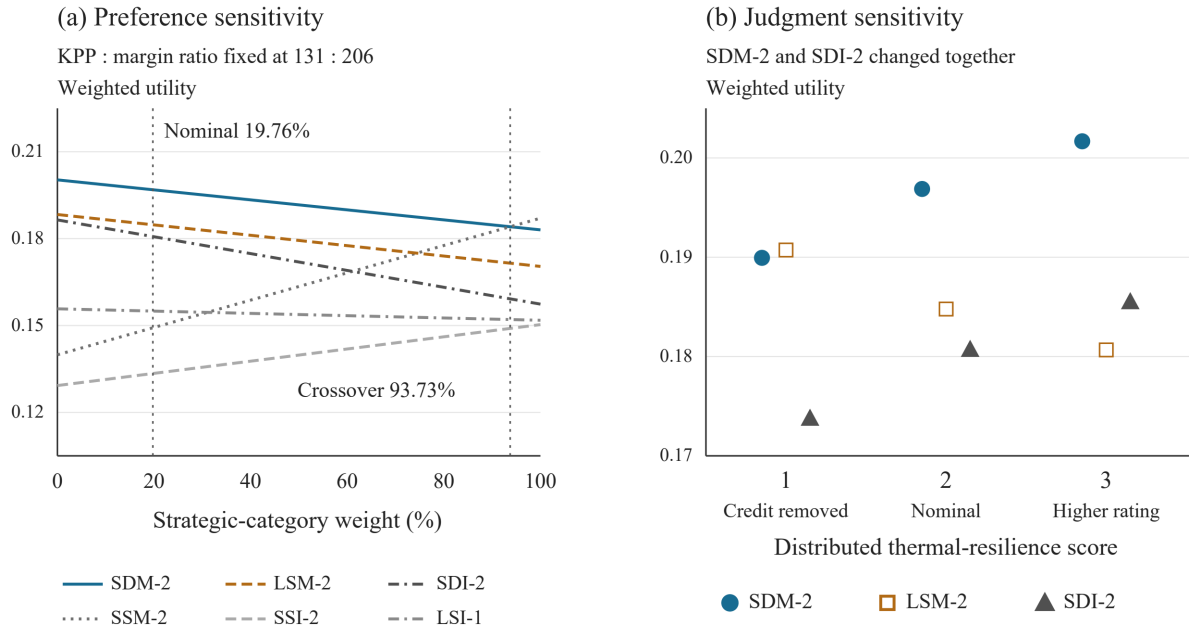

Fig. 7. Two bounded sensitivity experiments. (a) All six architectures under strategic-weight variation, with the KPP-to-margin ratio fixed. (b) The three leading contenders at nominal category weights when both distributed thermal-resilience scores are set to 1, 2, or 3; all six alternatives enter every calculation. Score 1 removes the fleet credit, 2 is the nominal assessment, and 3 is a hypothetical higher rubric rating. Utilities precede maximum-score scaling and are relative to the evaluated set. Discrete points do not represent measured survival probabilities or technology gains.

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

ture evaluation. It demonstrates inspection of decision assumptions, while measured designer-learning gains and a detailed technology-evolution analysis remain subjects for further work.

7. Conclusions

This paper presents a four-phase decision-support framework connecting strategic technology gaps, mission-derived operational capabilities, architecture comparison, and technology priorities in a New Space service context. Its contribution is the explicit organization of the interfaces between these decisions: strategic assessment, mission and capability artifacts, architecture evidence, and technology contributions remain inspectable alongside the engineering judgments that connect them.

The lunar SRP application identifies 16 gaps, 11 missions, 12 OCs, and 16 evaluation criteria. Twenty-four morphological combinations are reduced by context screening to 15 concepts; six representative architectures are then compared. Under the adopted baseline allocations and qualitative assessments, SA-SDM-2 ranks first, followed by SA-LSM-2 at 93.86 and SA-SDI-2 at 91.79. Distinct mission–OC coverage weighted by mission relevance gives distributed-system management the highest relative TA contribution for the nominally preferred concept, followed by onboard software and avionics.

The bounded sensitivity demonstration shows that SDM remains preferred throughout the tested KPP and margin weight paths and until strategic weight reaches 93.73%. With the nominal stored-energy penalty retained, removing the distributed thermal-resilience credit makes LSM narrowly preferred. Removing both adjustments leaves SDM first. This interaction identifies fleet energy and thermal assumptions that require joint engineering assessment, even when the ranking is stable under category-weight changes.

By making mission needs, design assumptions, and their consequences explicit, the framework supports better-informed early decisions and targeted reassessment. Further work should elicit criterion-specific preferences and value scales, assess alternative normalizations, couple capability-evolution scenarios to physical and cost models, and evaluate the workflow with other systems and design teams. The immediate engineering priority is to connect fleet energy reserves, survival strategies, and retained mission capability to physically supported scoring criteria.

Acknowledgments

The authors are pleased to express their gratitude to the STARlab team at Politecnico di Torino for their contributions to the development of the methodology and its application to the lunar SRP case study.

References

[1] PwC. *Lunar Market Assessment. Building the Lunar Economy: Sectorial Forecasts and Market Opportunity*. PricewaterhouseCoopers, Jan. 13, 2026.

[2] P. B. Niles and R. Stephan. “The Commercial Lunar Payload Services Initiative”. In: *2025 IEEE Aerospace Conference*. 2025, pp. 1–7. DOI: 10.1109/AERO63441.2025.11068423.

[3] European Space Agency. *ESA to Invite Companies to Connect with the Moon*. Dec. 21, 2022. URL: https://www.esa.int/Applications/Connectivity_and_Secure_Communications/ESA_to_invite_companies_to_connect_with_the_Moon (accessed 14.09.2026).

[4] L. Caraccio, A. Dellacasa, N. Di Gruttola Giardino, L. M. Festa, T. Galliano, G. Gorgerino, A. Meloni, F. Mustich, F. Stesina, and E. Vacchetto. “ARDITO, a Modular Technology Demonstrator for Robotic Planetary Surface Exploration and Operational Support: An Overview”. In: *IEEE Aerospace and Electronic Systems Magazine* 40.8 (2025), pp. 94–105. DOI: 10.1109/MAES.2025.3564224.

[5] NASA Jet Propulsion Laboratory. *CADRE: Cooperative Autonomous Distributed Robotic Exploration*. URL: https://www.jpl.nasa.gov/missions/cadre/ (accessed 14.09.2026).

[6] International Space Exploration Coordination Group (ISECG). *Global Exploration Roadmap 2024*. International Space Exploration Coordination Group (ISECG), 2024. URL: https://www.globalspaceexploration.org/?page_id=1371 (accessed 14.09.2026).

[7] National Aeronautics and Space Administration. *NASA Systems Engineering Handbook*. NASA Special Publication NASA/SP-2016-6105 Rev. 2. National Aeronautics and Space Administration, 2016. URL: https://ntrs.nasa.gov/api/citations/20170001761/downloads/20170001761.pdf (accessed 14.09.2026).

[8] J. M. Lafleur, J. L. Sharma, and J. Apa. “From Mission Objectives to Design: An Efficient Framework for Downselection in Robotic Space Exploration”. In: *AIAA SPACE 2007 Conference & Exposition*. AIAA 2007-6258. Long Beach, California, 2007.

[9] A. M. Ross, D. E. Hastings, J. M. Warmkessel, and N. P. Diller. “Multi-Attribute Tradespace Exploration as Front End for Effective Space System Design”. In: *Journal of Spacecraft and Rockets* 41.1 (2004), pp. 20–28. DOI: 10.2514/1.9204.

[10] M. Mosleh, K. Dalili, and B. Heydari. “Distributed or Monolithic? A Computational Architecture Decision Framework”. In: *IEEE Systems Journal*

Copyright © 2026 by the International Astronautical Federation (IAF). All rights reserved.

12.1 (2018), pp. 125–136. DOI: 10.1109/JSYST.2016.2594290.

[11] H. Chen, M. Ornik, and K. Ho. “Space Exploration Architecture and Design Framework for Commercialization”. In: *Journal of Spacecraft and Rockets* 59.2 (2022). DOI: 10.2514/1.A35077.

[12] T. L. Saaty. “Decision making with the Analytic Hierarchy Process”. In: *International Journal of Services Sciences* 1.1 (2008), pp. 83–98. DOI: 10.1504/IJSSCI.2008.017590.

[13] J. Figueira, S. Greco, and M. Ehrgott, eds. *Multiple Criteria Decision Analysis: State of the Art Surveys*. Vol. 78. International Series in Operations Research & Management Science. New York: Springer, 2005. DOI: 10.1007/b100605.

[14] G. Kalácska, G. Barkó, H. Shegawu, Á. Kalácska, L. Zsidai, R. Keresztes, and Z. Károly. “The Abrasive Effect of Moon and Mars Regolith Simulants on Stainless Steel Rotating Shaft and Polytetrafluoroethylene Sealing Material Pairs”. In: *Materials* 17.17, 4240 (2024). DOI: 10.3390/ma17174240.

[15] J. R. Hauser and D. Clausing. “The House of Quality”. In: *Harvard Business Review* 66.3 (1988), pp. 63–73.

[16] T. Ritchey. “General Morphological Analysis as a Basic Scientific Modelling Method”. In: *Technological Forecasting and Social Change* 126 (2018), pp. 81–91. DOI: 10.1016/j.techfore.2017.05.027.

[17] European Space Agency. *Argonaut: Europe’s Lunar Lander Programme*. 2025. URL: https://www.esa.int/Science_Exploration/Human_and_Robotic_Exploration/Exploration/Argonaut_Europe_s_lunar_lander_programme (accessed 14.09.2026).

[18] M. Shirley and E. Balaban. *An Overview of Mission Planning for the VIPER Rover*. SpaceOps 2022 Mission Planning Workshop. June 2, 2022. URL: https://ntrs.nasa.gov/citations/20220008301 (accessed 07.09.2026).