

Final missions: build a complete and resilient Mars architecture

English edition R77. This PDF preserves the instructional article, calculations, exercises, limits, decisions and available figures for the module.

Release source: R77-20260912.

Final missions: build a complete and resilient Mars architecture

Download this course as PDF

07 · Earth → Mars09 · Spacecraft11 · EDL12 · ECLSS14 · Reliability15 · Operations16 · Final missions

PERMANENT MARS SETTLEMENT

THE 6 PILLARS OF A SUSTAINABLE OUTPOST



No single pillar is enough: the settlement fails if one of them breaks for too long.



DELTA-SIERRA • MARS

The capstone does not judge one subsystem in isolation: transport, life support, power, habitat, industry, people and governance must remain coherent when failure or delay deforms the nominal plan.

The previous fifteen modules teach calculations, architecture reading and the limits of overly simple reasoning. Module 16 introduces almost no new concept. Instead, it forces them to connect. A propulsion equation meets a mass budget; a water-consumption estimate meets an emergency reserve; a communications outage meets local authority; and an attractive architecture has to show that it remains operable when reality stops following the nominal scenario.

A final study does not need to imitate a NASA or SpaceX programme. It must be labelled as a Delta-Sierra educational architecture: dated assumptions, reproducible calculations, visible margins and a clear boundary between capabilities demonstrated today and engineering extrapolation.

1. The mission brief: start from functions, not from a favourite vehicle

Each team receives a campaign objective: pre-deploy cargo, transport a crew, land essential elements, remain on the surface long enough to cross a return window, then leave an infrastructure that is more robust than it was on arrival. The first task is to translate this objective into functions: power, air, water, food, thermal control, communications, mobility, maintenance, medical care, refuge, storage, local production and return.

Reference crew Choose a crew size and explain why habitat, medical capability, workload and rescue capacity can support it.

Campaign duration Choose transit, surface stay and return window, then calculate consequences for consumables.

Pre-deployed cargo Identify what must work before humans arrive and what can wait.

Return option Define the power, propellant, spares and decisions required to keep return credible.

2. Build the architecture as functional chains

A useful architecture reads from input to output: source, conversion, storage, distribution, use, measurement, maintenance and degraded mode. Water is not just a tank. It needs a source, extraction, purification, storage, distribution, recovery, quality analysis and a response if contamination removes one loop from service.

For every life-critical function, the final sheet should identify the nominal path, backup path, endurance of the backup, fault-detection method, any shared resource capable of defeating both paths, and the expected human or automatic action. This table becomes the backbone of the design review.

3. Close the budgets before discussing appearance

An architecture that does not close its budgets is not yet an architecture. At minimum it should track mass, energy, peak power, water, oxygen, food, habitable volume, data storage, communications capacity, spare parts and crew time. Values may be educational assumptions, but their origin and sensitivity must be explicit.

Minimum protected stock — separate base demand, emergency reserve and uncertainty

$$S_{\min} = q_{\text{avg}} \times t_{\text{u}} + S_{\text{emg}} + U$$

1 — Concrete question

How much stock should a mission protect before entering a period with no dependable resupply or replacement production?

2 — Intuition without symbols

First cover expected consumption for the unsupported period. Then add a deliberately protected emergency reserve and a separate allowance for uncertainty so neither is hidden inside an optimistic average.

3 — Quantities first

S_{min} minimum protected stock; **q_{avg}** verified average consumption rate; **t_u** unsupported duration; **S_{emg}** emergency reserve reserved by policy; **U** uncertainty allowance for demand, loss or duration uncertainty.

4 — Formula

$$S_{\min} = q_{\text{avg}} \times t_{\text{u}} + S_{\text{emg}} + U$$

5 — Read aloud

“Minimum stock equals average consumption times unsupported duration, plus emergency reserve, plus uncertainty allowance.”

6 — Symbols

S means stock; q denotes a rate; t denotes time. The subscripts avg, u and emg mean average, unsupported and emergency.

7 — Pronunciation

q is read “cue”; S is read “ess”. Subscripts identify the role of each quantity and are not multipliers.

8 — Units

If q_{avg} is units/day and t_u is days, their product is units. S_{emg} and U must therefore also be expressed in the same stock unit.

9 — Convention

Do not count the same protection twice. If an uncertainty fraction m is used, define $U = m \times q_{\text{avg}} \times t_{\text{u}}$. Keep emergency reserve separate when it protects a distinct contingency.

10 — Why these quantities are related

Rate times time gives expected base consumption. The two additions represent different protections: one against an explicit emergency policy and one against uncertainty in the planning basis.

11 — Assumptions

The average rate is representative over the stated interval, stock is actually accessible and qualified for the intended use, and no unmodelled degradation makes part of the inventory unusable.

12 — Unit check

$(\text{units/day}) \times \text{day} = \text{units}$; adding reserve and uncertainty is valid only because all terms are stocks expressed in the same unit.

13 — Numerical case

$$q_{\text{avg}} = 18 \text{ units/day}$$

$$t_u = 45 \text{ days}$$

$$\text{Base demand} = 18 \times 45 = 810 \text{ units}$$

$$\text{Uncertainty fraction } m = 0.15$$

$$U = 0.15 \times 810 = 121.5 \text{ units}$$

$$S_{\text{emg}} = 81 \text{ units}$$

$$S_{\text{min}} = 810 + 121.5 + 81 = 1,012.5 \text{ units}$$

14 — Why each operation is performed

The first multiplication closes nominal consumption for the unsupported interval. The uncertainty allowance is computed from that base rather than from the already protected total. The emergency reserve is then added as a separate policy-protected stock.

15 — Algebra check

If available stock S_{avail} and all other quantities are known, the maximum unsupported duration before crossing the protection threshold is $t_u = (S_{\text{avail}} - S_{\text{emg}} - U) / (q_{\text{avg}})$, provided the numerator is positive and U is defined consistently.

16 — Mental estimate

Eighteen units per day for roughly forty-five days is a little over 800 units. Adding about 200 units of combined protection should place the minimum a little above 1,000 units, consistent with 1,012.5.

17 — Interpretation

The planning criterion protects 1,012.5 units under the teaching assumptions. A stock below that value does not mean immediate loss of life; it means the approved margin has already been consumed and the mission must act before raw endurance becomes the only remaining protection.

18 — What the result does not prove

It does not prove consumption stays constant, the entire stock is accessible, the product remains qualified, or resupply/production recovery will occur on schedule.

19 — Limit and sensitivity cases

If q_{avg} rises, t_u lengthens, or qualification losses reduce accessible stock, the required protection rises. If $S_{\text{avail}} \leq S_{\text{emg}} + U$, there is no positive unsupported duration under this policy; the architecture is already inside protected reserves.

20 — Practice

Guided exercise. A consumable is used at 12 units/day for 30 unsupported days. Protect 60 emergency units and use a 10% uncertainty allowance on base demand. Find $S_{\min}$.

Guided correction — open after attempting the guided exercise

Detailed guided correction.

1. Base demand = $12 \times 30 = 360$ units.
2. Uncertainty allowance = $0.10 \times 360 = 36$ units.
3. Add emergency reserve: $360 + 36 + 60 = 456$ units.
4. Minimum protected stock is therefore **456 units**.

Autonomous exercise. Available qualified stock is 860 units, emergency reserve is 80 units, uncertainty allowance is fixed at 100 units for the reviewed interval, and average use is 18 units/day. What unsupported duration can be defended before protected stock is crossed?

Autonomous correction — open after attempting the exercise

One defensible worked solution.

1. Stock available for planned consumption = $860 - 80 - 100 = 680$ units.
2. Duration = $680 / 18 \approx 37.8$ days.
3. A 45-day unsupported campaign therefore fails this protection criterion unless consumption falls, qualified stock rises, or the reserve policy is deliberately changed by the proper authority.

21 — Mission decision

Use the protected-stock calculation as a gate, not a decorative margin. If verified qualified inventory falls below the criterion, declare HOLD on activities that consume the constrained resource and open an explicit recovery or campaign-shortening decision.

If a resource is used at average rate q for a duration t , the base requirement is $q \times t$. A mission adds losses, variability, refuge days and delay scenarios. The goal is not to find a magical “exact” number; it is to identify which assumption drives the design and how the architecture changes when that assumption moves.

4. Final mission A — cargo must make human arrival less risky

Cargo leaves first. It carries the elements whose absence would make crew arrival unacceptable: initial power, communications, unloading capability, habitat or refuge, reserves, critical spares and possibly local-resource equipment. The study must state what is tested before crew departure, which telemetry demonstrates readiness and which threshold would force the human launch to be delayed.

Then inject a failure: a cargo flight misses its window, a solar array is damaged or a handling rover is immobilised. The team must show whether the site can continue preparing for arrival with reduced capability or whether the campaign has to be reconfigured.

5. Final mission B — human transit is a moving habitat

The crew vehicle must be assessed as a temporary habitat: air, water, thermal control, food, sleep, exercise, hygiene, medical capability, radiation shelter, maintenance and access to equipment. Propulsion alone is not the architecture. A shorter trip may move mass into propulsion; a slower trajectory may increase consumables and exposure. The correct answer is the one that makes the whole system more robust for the chosen scenario.

EARTH-MARS TRANSPORT

The complete chain of a crewed mission

1



PRE-DEPLOY THE CARGO

Cargo flights are sent ahead to prepare for the mission arrival.

2



LAUNCH THE CREW

The crew launches and joins the cargo in Earth orbit.

3



TRANSIT WITH LIFE SUPPORT

The spacecraft travels to Mars with all life-support systems operational.

4



LAND & DEPLOY

Safe landing, deployment, and activation of habitats and equipment.

5



ENSURE RETURN & NEXT WAVE

Crew return and preparation for the next mission.

A continuous cycle to build a sustainable human presence on Mars.



DELTA - SIERRA • MARS

The integrated mission follows the entire logistics chain: cargo, crew, transit, arrival, surface operations, return and preparation for the next wave.

6. Final mission C — arrival, EDL and site activation

The scenario must connect landing dispersion with surface operations. Two vehicles several kilometres apart may both count as successful landings while making the campaign much harder. The design therefore defines localisation, mobility, towing, energy reserves and the time needed to bring life-critical functions online.

The capstone must also state what happens if the crew arrives while part of the site is still off-nominal. Is there an independent refuge? How long can it operate? Which tasks can be deferred? Which operations remain prohibited until verification is complete?

7. Final mission D — turn an initial base into a durable system

A base that survives one hundred days is not automatically a settlement. The long campaign should show how maintenance, inventory, local production and training progressively reduce dependencies. A locally made part is useful only if material, geometry and quality are controlled. An extracted resource is useful only if it is purified, stored and distributed.

Sustainability therefore appears in loops: recycle more, repair more part families, manufacture selected consumables, increase refuge capacity, document procedures and transfer competence to the next team.

8. Long campaign: 500 sols without a perfect scenario

The operations plan should be long enough for ageing to matter. Filters clog, seals age, batteries lose capacity, dust accumulates, spares are consumed and people fatigue. The schedule must reserve time for inspection, exercise, care, training and lessons learned.

Do not fill every hour. An architecture optimised to full utilisation often assumes nothing unexpected will happen. Time margin is a resilience resource just like spare parts.

9. Injected failures: the mission is not validated until it has been disturbed

- **Logistics delay:** the next cargo arrives months later.
- **Power:** a dust event strongly reduces solar production.
- **ECLSS:** one oxygen-production train is unavailable for several days.
- **Communications:** Earth is temporarily unavailable or too slow for immediate decisions.
- **Mobility:** a pressurised rover is immobilised far from the site.
- **Medical:** one crewmember can no longer perform a primary role.

For each failure, the final dossier explains detection, decision authority, loads shed, available endurance, recovery criteria and the effect on the rest of the campaign. If the answer is “send a replacement from Earth,” it must include launch-window timing and the real delay before the part can become available.

10. Decision gates: know when to say “we do not launch”

Gate 1 — cargo departure: budgets closed, interfaces frozen, deployment capability verified.

Gate 2 — crew departure: critical site functions confirmed, reserves verified, abort criteria explicit.

Gate 3 — commit to surface: weather, EDL, navigation, communications and refuge are compatible with arrival.

Gate 4 — extend the campaign: stocks, health, power, water, maintenance and return remain within accepted limits.

A strong dossier is not trying to prove that the mission must fly. It makes visible the condition that requires delay or abort. The ability to say “no” is a mark of engineering maturity.

11. Final quantitative exercise — connect margin to a decision

For the exercise only, suppose a critical reserve is consumed at 18 units per day, the campaign must survive 45 days without outside support, the team protects an 81-unit emergency reserve, and uncertainty on the base demand is represented by a 15% allowance. The planning need is:

The protected-stock mini-lesson calculates the forty-five-sol base need step by step and obtains 810 units.

It then separates uncertainty allowance from emergency reserve and shows that the two protections total 202.5 units.

The same mini-lesson combines base need and the two protection layers to obtain a protected criterion of 1,012.5 units.

If measured qualified stock after a failure falls to 860 units, the question is not “can we survive today?” but “what decision do we make now so that the protected-stock gate is restored before a limit is crossed?” The team may reduce consumption, restore production, cancel an activity, change the campaign duration or use protected reserve only through the declared authority path.

Reasoned correction

The protected criterion is 1,012.5 units, so 860 qualified units are 152.5 units below the reviewed threshold. That does not imply immediate depletion; it means emergency and uncertainty protection are no longer intact. The board must either reduce demand, restore qualified stock, shorten the unsupported interval or consciously change the reserve policy. Raw endurance alone is not the approval criterion.

12. Review package: what the team must deliver

1. a mission-and-assumptions page;
2. a functional architecture and interface map;
3. main budgets with units, margins and sources;
4. a cargo-crew-surface-return timeline;
5. degraded modes and emergency stocks;
6. at least six injected failures;
7. a decision-and-authority matrix;
8. launch, continue, retreat and abort criteria;
9. a register of demonstrated, extrapolated and prospective elements;
10. a conclusion naming the three largest residual risks.

13. Evaluate the architecture without rewarding spectacle

The final assessment should reward coherence, traceability and the ability to acknowledge limits. A less spectacular architecture with explicit budgets, backups and decisions is stronger than an impressive concept built on hidden assumptions. The most important criteria are closure of life-critical functions, common-cause control, repairability, margin discipline and the boundary between fact and prospect.

14. The last skill: hand the system to the next team

A durable human mission does not end at return. It leaves data, procedures, qualified parts, known errors and a base that is easier for the next crew to understand. The final documentation therefore explains not only what was built, but what changed during the campaign and why.

At this point the student does not “know Mars” in an encyclopaedic sense. They can do something more useful: state an assumption, quantify it, connect it to other systems, search for what can break it, and explain a decision in a form that another person can verify.

Reference architecture dossier — close the mission before approving it

This integrative extension turns mission architecture into a reviewable chain of budgets, dependencies, degraded states, recovery gates and return-preservation decisions that a learner can defend under challenge.

1. Build an assumption register that can be challenged

The final mission should begin with an assumption register rather than a polished diagram. Every major input is tagged as requirement, measured value, heritage value, model output or training assumption. The register records units, source, uncertainty and the design decisions that depend on it. If a value changes, reviewers can see which budgets must be reopened. This prevents hidden assumptions from becoming structural features of the mission simply because they were copied from an earlier worksheet.

2. Close water, atmosphere and food as time-dependent inventories

For each consumable, separate initial stock, production or recovery rate, unavoidable losses, contingency stock and protected reserve. A daily average is useful for first sizing but must be tested against maintenance outages and demand peaks. Food differs from water because most mass is not regenerated on short timescales; oxygen differs because compressed or cryogenic storage can provide a rapid emergency buffer while production equipment recovers. The mission dossier should state exactly how long each critical inventory can bridge the loss of its regeneration function.

3. Close power with load classes and restart logic

List continuous, cyclic and contingency loads separately. Critical life-support, communications and thermal loads should have a protected supply path and a load-shedding order. The energy budget must include duration, not only peak kilowatts. A restart plan matters because several systems may demand high power simultaneously after an outage. Recovery sequencing should prevent the attempt to restart everything from collapsing the power system a second time.

4. Treat crew time as a finite mission resource

Every maintenance action, sample campaign, EVA, inventory check and manual degraded-mode task consumes crew time. Summing planned hours by skill category exposes hidden staffing problems. A mission that closes in mass and power can still be infeasible if the same two specialists are scheduled for overlapping critical tasks. The dossier should therefore include a workload budget, cross-qualification plan and margin for anomaly response.

5. Prove rescue and safe-haven logic

A rescue concept is credible only when time, distance, consumables, communications and medical capability are closed together. A rover range quoted by the manufacturer is not automatically a rescue radius because return energy, terrain, thermal conditions and degraded operation consume margin. Safe havens require their own atmosphere, power, water and communications assumptions. The architecture should show what happens if the primary habitat cannot be re-entered.

6. Inject failures before declaring success

Run the mission timeline with failures deliberately inserted: delayed cargo, partial power loss, water-quality uncertainty, communications outage, EVA injury and a critical spare shortage are examples. The purpose is not to invent drama but to test whether buffers and authority rules work under combined stress. Record the state before the event, detection path, immediate protective action, recovery path, remaining margin and criteria for continuing the mission.

7. Use GO / HOLD / NO-GO gates with measurable evidence

A launch or surface campaign should not proceed because the calendar says it should. Define gates whose evidence is measurable: verified inventory above a protected threshold, required redundancy available, unresolved hazards below an agreed severity, qualified operators present and critical software/configuration baselines frozen. HOLD means evidence is incomplete or a recoverable condition must be corrected. NO-GO means a requirement is not met and the architecture must change.

8. Deliver a review package that another team can inherit

The final product is not only a presentation. It includes architecture diagrams, budgets with units, source register, configuration baseline, risk register, maintenance concept, failure-injection results, decision logs and open issues. A new team should be able to reconstruct why the mission was judged acceptable without relying on the memory of the original designers. That handover quality is itself a measure of engineering maturity.

Progressive mastery drills — eight linked checks

Drill 1 — Assumption register

Classify each major input as requirement, measurement, heritage value, model output or training assumption.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Drill 2 — Water closure

Build initial stock, recovery, loss, reserve and outage-duration lines in one budget.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Drill 3 — Atmosphere closure

Separate oxygen production, stored reserve, leakage, EVA losses and emergency supply.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Drill 4 — Food closure

Distinguish stored calories, locally produced fraction, crop failure reserve and resupply dependence.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Drill 5 — Power closure

Create critical, essential and deferrable load classes and a restart sequence after blackout.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Drill 6 — Crew-time closure

Sum maintenance and operations by skill category and identify overload during anomaly response.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Drill 7 — Mobility and rescue

Turn rover energy, terrain and return reserve into a rescue radius rather than using nominal range.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Drill 8 — Failure-injection review

Run one combined power-plus-water event through detection, protection, recovery, residual margin and GO/HOLD decision.

Expected reasoning. Show the physical meaning, units or evidence path, state at least one assumption, and explain what operational decision would change if the result or evidence were different.

Integrated exercise — Integrated review exercise — defend a HOLD decision

In a training architecture, cargo has delivered the nominal water reserve and the surface power system passes a normal-load test. However, one redundant water-quality sensor remains uncalibrated, a critical pump spare has not arrived, and the crew workload model shows no margin during the first two weeks. Prepare a short review-board decision using GO, HOLD or NO-GO. Identify the evidence required to change the decision.

Reasoned solution. HOLD is defensible because several recoverable closure items remain unresolved and they interact with life-support resilience. The board should require a validated water-quality measurement path, a credible pump-failure recovery plan or delivered spare, and a revised workload/shift plan that preserves anomaly margin. GO becomes reasonable only after those closure items are evidenced; NO-GO would be appropriate if they cannot be resolved within the mission architecture or launch window.

Primary-source bridge for this extension. Official source 1 Official source 2 Official source 3. These links are evidence anchors already present in the French reference module; values and limits must still be checked against the source version applicable to the mission.

Mission-resource equations — turn inventories and mobility into protected horizons

The capstone repeatedly uses water, mobility and reserve calculations. The following mini-lessons make two of the most important relationships explicit so a review board can reproduce them without relying on prose arithmetic.

Campaign water make-up — convert verified net loss into inventory mass

$$M_{\text{makeup}} = q_{\text{net}} \times N \times t$$

1 — Concrete question

How much net water make-up mass is required for a crew over a specified campaign when recovery losses are already represented by a verified per-person net make-up rate?

2 — Intuition without symbols

Multiply the verified net daily loss per person by crew size and campaign duration. This is a make-up inventory calculation, not a full water-system model.

3 — Quantities first

M_{makeup} is campaign make-up mass; q_{net} is verified net make-up per person per sol; N is crew count; t is campaign duration in sols.

4 — Formula

$$M_{\text{makeup}} = q_{\text{net}} \times N \times t$$

5 — Read aloud

"M make-up equals q net times N times t."

6 — Symbols

M_makeup is mass to be supplied; q_net is net make-up rate; N is number of people; t is duration.

7 — Pronunciation

q is "cue"; N is "en"; t is "tee".

8 — Units

kg/(person·sol) × person × sol = kg.

9 — Convention

q_net must already correspond to the defined operating state and water-quality boundary. Do not silently mix potable, hygiene and process water.

10 — Why this operation

Each person incurs the net loss each sol. Crew size scales the daily mission loss; duration scales it over the campaign.

11 — Assumptions

The net rate is treated as constant over the interval; outages, leaks, startup losses and contingency reserves are handled separately.

12 — Unit check

Person and sol cancel, leaving kilograms.

13 — Numerical case

$q_{\text{net}} = 1.2 \text{ kg}/(\text{person} \cdot \text{sol})$.

$N = 6 \text{ people}$.

$t = 500 \text{ sols}$.

Daily crew make-up = $1.2 \times 6 = 7.2 \text{ kg/sol}$.

$M_{\text{makeup}} = 7.2 \times 500 = 3,600 \text{ kg}$.

14 — Why each operation

First scale the per-person loss to the whole crew, then accumulate that crew loss across campaign duration.

15 — Algebra check

Solving for endurance gives $t = M_{\text{available}} / (q_{\text{net}} N)$ when the same constant-rate assumptions hold.

16 — Mental estimate

About 7 kg/sol for about 500 sols is about 3,500 kg, so 3,600 kg is plausible.

17 — Interpretation

The teaching scenario needs 3.6 tonnes of net make-up over 500 sols before separately adding protected reserves or scenario-specific losses.

18 — What it does not prove

It does not prove the ECLSS can actually recover, qualify, store or distribute water, and it does not include startup, maintenance or emergency losses.

19 — Sensitivity or limit case

If q_{net} doubles during a degraded period, campaign consumption rises immediately. If q_{net} approaches zero, make-up demand falls but verification of quality and redundancy still matters.

20 — Practice

Guided exercise. For $N=4$, $q_{\text{net}}=1.5 \text{ kg}/(\text{person}\cdot\text{sol})$, $t=30$ sols, calculate make-up mass.

Guided correction — open after attempting the guided exercise

Detailed guided correction.

1. Daily crew make-up = $1.5 \times 4 = 6 \text{ kg/sol}$.
2. Campaign make-up = $6 \times 30 = 180 \text{ kg}$.
3. Keep any emergency buffer separate from this nominal make-up result.

Autonomous exercise. A six-person campaign has 600 kg of accessible potable reserve and $q_{\text{net}} = 2.0 \text{ kg}/(\text{person}\cdot\text{sol})$ after degradation. Estimate the constant-rate horizon and state two reasons not to treat it as guaranteed survival time.

Autonomous correction — open after attempting the exercise

One defensible worked solution.

1. Daily crew make-up = $2 \times 6 = 12 \text{ kg/sol}$.
2. Horizon = $600/12 = 50$ sols.
3. Real survival also depends on water quality, distribution, power, maintenance, medical demand and whether the 600 kg is actually qualified and accessible.

21 — Mission decision

Use the formula to expose inventory horizon, then protect separate contingency reserves and verify the recovery function. A mass total without quality and accessibility is not mission closure.

Protected rescue radius — preserve the return leg before leaving the habitat

$$R_{\text{safe}} = (R_{\text{nom}} \times f_{\text{usable}}) / 2$$

1 — Concrete question

Given a nominal rover range and a protected reserve fraction, what simple out-and-back radius should not be exceeded before terrain and environmental penalties are added?

2 — Intuition without symbols

A rescue rover must travel out and still return. First protect part of the advertised range, then divide the usable remainder between outbound and inbound travel.

3 — Quantities first

R_{safe} is simple safe radius; R_{nom} is nominal range under defined conditions; f_{usable} is the fraction allowed for planned travel after reserve protection.

4 — Formula

$$R_{\text{safe}} = (R_{\text{nom}} \times f_{\text{usable}}) / 2$$

5 — Read aloud

“ R_{safe} equals R_{nom} times f_{usable} , divided by two.”

6 — Symbols

R_{safe} is one-way planning radius; R_{nom} is nominal total range; f_{usable} is a dimensionless fraction between zero and one.

7 — Pronunciation

R is “ar”; f is “eff”.

8 — Units

Distance $\times$ dimensionless fraction $\div 2$ remains distance.

9 — Convention

This is symmetric out-and-back geometry. Terrain, detours, cold-battery effects, towing and rescue payload must reduce the result further when applicable.

10 — Why this operation

Reserve protection reduces usable total range. Dividing by two allocates that usable distance to an outbound and return leg.

11 — Assumptions

Same effective range outbound and inbound, no alternate refuge, no recharge and no terrain/ weather penalty in the first-order geometry.

12 — Unit check

$\text{km} \times 0.75 \div 2$ gives km.

13 — Numerical case

$R_{\text{nom}} = 60 \text{ km.}$

Protected reserve = 25%, so $f_{\text{usable}} = 0.75.$

Usable total travel = $60 \times 0.75 = 45 \text{ km.}$

$R_{\text{safe}} = 45 / 2 = 22.5 \text{ km.}$

14 — Why each operation

Convert the reserve policy to a usable fraction, apply it to nominal range, then split the remaining distance between going out and coming back.

15 — Algebra check

A target at distance d requires at least $R_{\text{nom}} \geq 2d/f_{\text{usable}}$ before adding terrain and contingency penalties.

16 — Mental estimate

Three quarters of 60 is 45; half is 22.5.

17 — Interpretation

Under the simplified assumptions, 22.5 km is the maximum symmetric planning radius before additional penalties.

18 — What it does not prove

It does not certify rescue success, battery temperature, route passability, towing performance, communications, crew time or safe-haven availability.

19 — Sensitivity or limit case

A larger protected reserve lowers radius. An asymmetric route or one-way safe haven requires a different geometry and should not be forced into this formula.

20 — Practice

Guided exercise. For $R_{\text{nom}}=80$ km with 30% protected reserve, compute the simple radius.

Guided correction — open after attempting the guided exercise

Detailed guided correction.

1. $f_{\text{usable}} = 0.70$.
2. Usable total = $80 \times 0.70 = 56$ km.
3. $R_{\text{safe}} = 56/2 = 28$ km.

Autonomous exercise. A casualty is 24 km away. The rover has 70 km nominal range, 20% protected reserve, and terrain is expected to add 15% distance to each leg. Does the simple nominal formula close the rescue?


Autonomous correction — open after attempting the exercise

One defensible worked solution.

1. Simple usable range = $70 \times 0.80 = 56$ km, giving 28 km radius before terrain.
2. Terrain-adjusted round trip is $2 \times 24 \times 1.15 = 55.2$ km.
3. Only 0.8 km remains inside the usable 56 km allowance, before towing/cold/weather effects.
4. Disposition should be HOLD pending a stronger margin, alternate support or revised route; the simple radius alone is not robust enough.

21 — Mission decision

Treat mobility radius as an option-preservation gate. If penalties and uncertainty consume the protected margin, shorten the operating envelope or add another rescue path before committing the crew.



Mars mission-operations scene with deterministic protected-function horizons. The photographic layer supplies operational context; the GO/HOLD/NO-GO bars and labels are exact review data.

Sources and pathways

- NASA — Human Spaceflight
- NASA — Space Technology Mission Directorate
- ESA — Human and Robotic Exploration
- Mars Bible — Earth-Mars transport
- Mars Bible — ECLSS and habitat
- Mars Bible — reliability and common causes

First Man capstone dossier — quantitative mission closure before approval

The final mission module should force every earlier discipline to meet in one review package. The numbers below are a deliberately simplified teaching scenario, not a NASA, ESA or commercial mission design. Their purpose is to make assumptions visible, close inventories, expose interfaces and train GO/HOLD/NO-GO reasoning that another team can reproduce.

Reference scenario and assumption register

Scenario: six people conduct a 500-sol surface campaign after two precursor cargo missions. The precursor phase must already have demonstrated minimum power, communications, pressurizable shelter, consumable reserve and rescue mobility. The surface architecture includes a main habitat, an independently supportable refuge, two rovers able to support rescue, multiple power paths and enough protected inventory to enter degraded modes without immediately consuming the mission's last option.

Every number in this dossier is tagged as a teaching assumption. In a real design, it would be replaced by an approved requirement, measured test result, validated model or sourced design input. The assumption register is therefore part of the deliverable, not a hidden spreadsheet tab.

Close water, oxygen and food as time-dependent inventories

The reference campaign uses a verified net water make-up rate for six people over 500 sols. The complete inventory arithmetic now lives in the dedicated campaign water make-up mini-lesson, while the separate protected buffer remains a distinct contingency decision rather than being hidden inside the nominal campaign total.

For oxygen, use a teaching net basis of 0.9 kg/person/sol only to demonstrate inventory logic. Six people over 500 sols gives 2,700 kg on that simplified basis. Real design would separately model metabolic demand, leakage, EVA, pressurization, carbon-dioxide removal, gas buffering and production. For dry food at 0.75 kg/person/sol, the teaching mass is 2,250 kg before packaging, contingency and local production effects.

Mission closure matrix

Functions versus evidence, margin, degraded mode and decision gate.

Function	Evidence	Margin	Degraded mode	Gate
Water	test / measurement	protected reserve	verified fallback	GO / HOLD
Atmosphere	test / measurement	protected reserve	verified fallback	GO / HOLD
Power	test / measurement	protected reserve	verified fallback	GO / HOLD
Mobility	test / measurement	protected reserve	verified fallback	GO / HOLD
Crew time	test / measurement	protected reserve	verified fallback	GO / HOLD

Delta-Sierra pedagogical synthesis · schematic, not to scale

Functions versus evidence, margin, degraded mode and decision gate. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Close power as both power and energy

Assume an average active load of 55 kW and a critical degraded load of 35 kW. The worked power-budget development uses the declared Martian-sol duration and produces about 1,356.3 kilowatt-hours per sol for the average-load training scenario. A battery may be rated at 900 kWh nominal, but protected reserve and conversion losses make the usable emergency energy lower.

The architecture review therefore needs load classes, source availability, restart sequencing and the time at which a protected energy reserve would be crossed. A battery endurance number without the assumed critical load is meaningless.


Rescue mobility is a geometry and option-preservation problem

The reference mobility case starts from a nominal rover range, protects an operational reserve and preserves an out-and-back return leg. The complete first-order geometry now lives in the dedicated protected rescue-radius mini-lesson; terrain, detours, weather, towing and battery temperature remain explicit penalties rather than being hidden inside the nominal radius.

The rescue decision should confirm crew condition, location, remaining life-support endurance, the rescue rover's actual energy state, path risk and alternative refuges before dispatch. 'Inside range' is not the same as 'safe to launch the rescue'.

Injected-failure timeline

Six failures across a 500-sol campaign with recovery gates and protected reserves.



Delta-Sierra pedagogical synthesis · schematic, not to scale

Six failures across a 500-sol campaign with recovery gates and protected reserves. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Crew time is a finite mission resource

Six people with eight schedulable work hours provide 48 person-hours per sol, but reserving 20% for coordination, anomalies, documentation and recovery leaves only 38.4 person-hours for normal planning. Heavy maintenance, major EVA science, cargo inventory and emergency training cannot all be scheduled as if labour were free.

When an anomaly consumes 15 unplanned person-hours, record what was deferred and what risk that deferral creates. Recovery can generate a maintenance debt, inspection debt, sleep debt or documentation debt that survives long after the alarm disappears.

Six injected failures test whether the architecture is actually resilient

Failure 1: primary power source unavailable for 18 h. Shift to the 35 kW critical profile, compute storage endurance and place nonessential activities on HOLD. Failure 2: water recovery degrades for ten sols, increasing net makeup demand; calculate how much protected stock is consumed. Failure 3: primary oxygen production is unavailable for three sols; verify gas reserve and atmosphere-quality control independently.

Failure 4: one rover is immobilized 18 km away; preserve rescue-vehicle reserve and alternative safe-haven options. Failure 5: high-bandwidth Earth communication is unavailable for 48 h; local procedures must continue without real-time terrestrial permission. Failure 6: one third of habitable volume is isolated; confirm that remaining atmosphere, sleeping, sanitation, thermal and human-factors limits are still acceptable.

GO, HOLD and NO-GO are evidence states, not moods

NASA's NASA — Moon to Mars Architecture provides modern architecture context, while NASA — Environmental Control and Life Support System grounds life-support functions. A GO means the required function and margin are demonstrated for the next phase. HOLD means the system is presently safe but evidence or margin is insufficient to proceed. NO-GO means a required function, rescue path or minimum margin is not demonstrated.

A review board should be able to point to the exact evidence that changes one state into another. 'We think the repair will work' is not release evidence. A verified test, stable trend, independent check or restored redundant path can be.

The handover package is part of mission success

Human-spaceflight and technology-development programs, such as the work represented by NASA — Human Spaceflight and NASA — Space Technology Mission Directorate, depend on configuration control and knowledge transfer. The final student deliverable should include requirements, assumption register, interfaces, budgets, hazards, test evidence, open waivers, spares, software/configuration state and a chronological anomaly log.

The capstone is therefore not passed because one architecture has impressive performance. It is passed when another team can reconstruct why the architecture was accepted, which margins remain, which failures were demonstrated, and what conditions would force a HOLD or NO-GO later.

Master closure table — one page that exposes every protected margin

Budget	Teaching nominal	Protected margin / reserve	Release question
Water makeup	7.2 kg/sol for six people	The stated water buffer corresponds to thirty sols at the declared nominal net loss in this training scenario.	Is the reserve independent of the failed recovery branch?
Oxygen net basis	5.4 kg/sol	Example 45 kg emergency stock	Does another function share the same stock?
Dry food	4.5 kg/sol	Campaign + contingency policy	What remains after schedule slip?
Average energy	The worked power budget produces about 1,356.3 kilowatt-hours per sol for the declared training scenario.	Protected battery reserve	Which source restores generation before endurance expires?
Critical power	35 kW	18.9 h teaching battery endurance	Are all life-critical loads really inside the profile?
Rescue mobility	22.5 km simple safe radius	25% nominal range protected	Do terrain and weather preserve this radius?
Crew time	48 person-h/sol gross	20% workload reserve	Which tasks are deliberately abandoned during an anomaly?

The closure table is intentionally compact. It is not the design model; it is the review index that tells the board where the detailed evidence lives. Every row should link to the calculation, test or procedure that justifies the nominal value and reserve. If one row contains an assumption that has not been demonstrated, the table should mark that status explicitly rather than presenting the number as settled fact.

Six injected failures — chronological reasoning rather than independent homework

Power loss, hour 0. Confirm the source is actually unavailable, isolate nonessential loads and verify the 35 kW critical profile. The calculated teaching endurance is about 18.9 h before the protected battery boundary. If the best credible repair path is 14 h, the raw time margin is about 4.9 h, but the board should still examine uncertainty in repair duration, battery temperature and actual load.

Water degradation, sol 120. Net makeup demand rises from 1.2 to 4.0 kg/person/sol for ten sols. The worked incident arithmetic produces 168 kg of extra water demand over the ten-sol degraded interval. The incident consumes much of the 216 kg thirty-sol teaching buffer. The correct response is not simply to deduct stock; it is to establish a new horizon, restrict nonessential uses and avoid spending the remaining reserve on convenience.

Oxygen-production outage, sol 205. At the teaching net basis, the three-sol requirement is 16.2 kg. Against an example 45 kg emergency inventory, arithmetic reserve after three sols is 28.8 kg if no other use draws from it. The crew still needs independent carbon-dioxide control and atmosphere verification. Oxygen mass alone is not a habitable-atmosphere certificate.

Rover immobilization, sol 278. One rover stops 18 km from the habitat. The simple safe radius is 22.5 km, leaving only 4.5 km geometric margin. Before sending the second rover, verify route length, detours, battery state, suit or cabin endurance, weather, communications and the possibility of using an alternate refuge. Rescue uses the same mobility reserve needed to bring the rescuers home.

Earth-communication loss, sol 341. High-bandwidth contact is unavailable for 48 h. Vital operations continue under pre-authorized local procedures. The crew logs decisions, compresses essential telemetry and distinguishes actions that merely require later notification from those that were designed to require prior external approval. If a vital function cannot proceed without real-time Earth authority, the mission architecture has not respected Mars communication delay and outage reality.

Habitable-volume isolation, sol 417. One third of the pressurized volume is taken out of service. The remaining habitat must support six people with acceptable atmosphere processing, sleeping, sanitation, thermal control, emergency egress and workload. The event tests human factors as much as hardware. A technically breathable volume can still be operationally unsustainable if crowding, noise, access or sleep degradation persists too long.

GO / HOLD / NO-GO matrix — decisions tied to evidence

State	Meaning	Teaching example
GO	Required function and margin demonstrated	Critical power path verified with endurance beyond repair horizon plus protected margin
HOLD	Present state safe but next step lacks evidence or margin	Water inventory adequate, but degraded-loss rate not yet stabilized
NO-GO	Required function, rescue path or minimum margin not demonstrated	No credible source can restore power before protected reserve is crossed
SAFE / RETURN MODE	Abandon nominal objective to preserve crew and reversible options	Cancel EVA, ration loads, occupy refuge or prepare return depending on mission phase

The student is not graded for reproducing these exact teaching numbers. A different architecture may choose different crew size, loads, storage, reserves or mission duration. It passes only if assumptions are explicit, units close, interfaces are visible, margins are protected and the decision logic remains coherent when the numbers change.

Final board exercise — challenge the architecture instead of presenting it

Prepare a twenty-minute review in which another team is instructed to attack the design. Give them the assumption register and closure table first, not a promotional rendering. They should choose one life-critical function, trace its power, data, thermal, human and spare-part dependencies, then inject a failure that removes one shared dependency. Your task is to show where the architecture detects the failure, how long the crew remains safe, which action preserves options and what evidence permits recovery.

The review is successful when the attacking team can reconstruct the logic without asking what an unlabeled number means. Any assumption that exists only in the presenter's memory is a configuration-control defect. Any reserve that is consumed by two budgets at once is double-counted. Any GO criterion that cannot be measured is not a criterion. The final competence is therefore the ability to make mission reasoning portable between teams.

Review drills — move from explanation to operational judgement

1. **Budget closure.** Build a one-page table for water, oxygen, food, energy, mobility and crew time with nominal need, protected reserve and release criterion.
2. **Failure injection.** For each of the six failures, identify the first evidence to collect and one action that preserves options.
3. **Review-board drill.** Write one measurable GO, HOLD and NO-GO statement for critical power.
4. **Handover drill.** List ten items the next mission team needs in order to reproduce the acceptance decision.

Usable battery endurance for a critical load

$$t_{\text{bat}} = E_{\text{nom}} \times (1 - r_{\text{reserve}}) \times \eta / P_{\text{crit}}$$

1 — Concrete question

How long can the nominal battery support the critical load after protected reserve and delivery efficiency are accounted for?

2 — Intuition without symbols

Start with nominal stored energy, remove the fraction that policy says must remain protected, account for the fraction of remaining energy actually delivered, then divide usable energy by the critical power draw.

3 — Quantities first

E_{nom} nominal battery energy; r_{reserve} protected fraction; η delivery efficiency; P_{crit} critical load; t_{bat} endurance.

4 — Formula

$$t_{\text{bat}} = E_{\text{nom}} \times (1 - r_{\text{reserve}}) \times \eta / P_{\text{crit}}$$

5 — Read aloud

"t battery equals E nominal times one minus r reserve times eta, divided by P critical."

6 — Symbols

t_{bat} endurance; E_{nom} nominal stored energy; r_{reserve} protected fraction; η efficiency; P_{crit} critical power.

7 — Pronunciation

η is read "eta"; r is read "are".

8 — Units

kWh $\times$ dimensionless $\times$ dimensionless $\div$ kW = h.

9 — Convention

Enter reserve and efficiency as fractions, not percentage numbers: 20% = 0.20 and 92% = 0.92.

10 — Why this operation

Nominal energy is reduced by the protected reserve and delivery losses; energy divided by power gives time.

11 — Assumptions

The model assumes constant critical load, usable capacity equal to the stated nominal basis, and no temperature, aging or power-limit derating beyond η and reserve.

12 — Unit check

kWh/kW = h.

13 — Numerical case

Nominal energy: E_{nom} = 900 kWh.

Protected reserve: r_{reserve} = 0.20.

Energy after reserve = 900 $\times$ 0.80 = 720 kWh.

Delivery efficiency: η = 0.92.

Delivered usable energy = 720 $\times$ 0.92 = 662.4 kWh.

Critical load: P_{crit} = 35 kW.

t_{bat} = 662.4 / 35 = 18.93 h.

14 — Why each operation

Subtract the reserve fraction from one to obtain the spendable share, multiply by efficiency to estimate delivered energy, then divide by the rate at which the critical load consumes energy.

15 — Algebra check

Required nominal energy for a target endurance is $E_{\text{nom}} = t_{\text{bat}} P_{\text{crit}} / [(1-r_{\text{reserve}})\eta]$.

16 — Mental estimate

About 660 kWh usable at about 33–35 kW should last around 19–20 h, matching the exact result.

17 — Interpretation

Under the teaching assumptions, the battery can support the 35 kW critical profile for about 18.9 h before the protected reserve boundary is reached.

18 — What it does not prove

It does not prove the battery can provide 35 kW continuously at the relevant temperature, nor that all critical loads are actually included.

19 — Sensitivity or limit case

If critical load rises to 45 kW, endurance falls to $662.4/45 \approx 14.72$ h; workload growth therefore consumes repair time quickly.

20 — Practice

Guided exercise. With $E_{\text{nom}} = 1,200$ kWh, reserve = 25%, $\eta = 0.90$ and $P_{\text{crit}} = 40$ kW, find endurance.

Guided correction — open after attempting the guided exercise

Detailed guided correction.

1. Spendable nominal share = $1,200 \times 0.75 = 900$ kWh.
2. Delivered usable energy = $900 \times 0.90 = 810$ kWh.
3. Endurance = $810 / 40 = 20.25$ h.

Autonomous exercise. A repair team estimates 16 h to restore generation. The battery is 900 kWh nominal with 20% protected reserve and 92% efficiency. What is the maximum constant critical load that still leaves a 2 h time margin beyond the estimate?

Autonomous correction — open after attempting the exercise

One defensible worked solution.

1. Required endurance = $16 + 2 = 18$ h.
2. Delivered usable energy = $900 \times 0.80 \times 0.92 = 662.4$ kWh.
3. Maximum critical load = $662.4 / 18 = 36.8$ kW.
4. If the verified critical profile exceeds 36.8 kW, shed additional nonessential load, improve repair confidence or provide another source before declaring GO.

21 — Mission decision

Tie battery endurance to repair-time evidence and a protected minimum margin. If the credible repair horizon exceeds demonstrated endurance, the state is HOLD or NO-GO rather than optimistic GO.

Final capstone review — can another team reproduce the mission decision?

- Is every mission-critical assumption visible in a register?
- Are water, atmosphere, food, energy, mobility and crew-time budgets all closed?
- Are reserves protected from double counting?
- Have injected failures crossed subsystem boundaries rather than remaining isolated?
- Do GO/HOLD/NO-GO criteria point to measurable evidence?
- Does the handover package preserve configuration, waivers and recovery debt?

The capstone is complete only when the architecture can be attacked constructively. Reviewers should be able to change a load, delay a repair, remove a shared dependency or invalidate an assumption and then watch the decision logic update without collapsing into improvisation. That is the First Man standard being pursued here: not memorizing a design, but learning to reason through one.

Primary-source bridges used in this closure dossier


- NASA — Moon to Mars Architecture
- NASA — Environmental Control and Life Support System
- NASA — Human Spaceflight
- NASA — Space Technology Mission Directorate

Closure rule. The capstone closes only when another team can reproduce the assumptions, budgets, failures, margins and GO/HOLD/NO-GO logic without relying on presenter memory.

Integrated mission closure board — defend a 500-sol architecture under cross-examination

A mission is not closed when every subsystem has a slide. It is closed when the interfaces, reserves, verification evidence, crew workload and recovery paths form one coherent argument. This capstone extension turns the existing budgets and six injected failures into a board-level dossier that another team could reproduce, challenge and inherit.

Integrated mission timeline



Precursor cargo, crew departure, transit, EDL, 500-sol operations and return are connected as evidence gates. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this dossier; not a mission-certified drawing.

Precursor cargo must prove functions before people depend on them

Primary-source bridge: NASA — Moon to Mars Architecture.

A cargo mission is not successful merely because it lands. The architecture needs a commissioning sequence that proves power, communications, environmental protection, mobility, stored consumables and any pre-deployed production. The evidence package must distinguish installed hardware from available hardware and available hardware from verified mission-ready capability.

The human departure gate should therefore reference explicit precursor evidence: what ran, for how long, under what environmental conditions, with which sensors, after which maintenance actions, and with what remaining margin. If a critical function is only inferred from telemetry or has never been exercised in the required configuration, the correct state may be HOLD even when the cargo vehicle itself is healthy.

Departure is a commitment gate, not a calendar date

Once the crew leaves Earth, some recovery options disappear and others become slower. A launch window therefore has to intersect with technical readiness, medical readiness, trajectory opportunities, surface assets, communication capability and return strategy. Schedule pressure is not evidence.

A board should be able to point to each protected margin and say what event would consume it. If the mission needs a last-minute waiver, the waiver must identify the lost protection and the compensating control. 'We are late' is never a compensating control.

Transit closes life support, medical capability and human performance together

The transit phase couples consumables, maintenance, exercise, sleep, radiation exposure, medical capability, communications delay and crew workload. A water reserve is not independent of the repair burden of the water loop; medical consumables are not independent of the crew's ability to diagnose and treat; exercise time competes with maintenance and science time.

The integrated dossier should therefore include not just masses but operational time: daily preventive maintenance, contingency maintenance, monitoring, training refresh and the workload created by degraded modes. A design that closes kilograms but requires impossible crew-hours does not close.

Arrival and EDL create a temporary peak-risk configuration

Primary-source bridge: NASA — Environmental Control and Life Support System.

Landing is not the end of EDL risk. After touchdown the crew may need to transition from vehicle survival mode to surface systems while communications, mobility and local power are still being confirmed. The first hours can therefore combine fatigue with incomplete infrastructure.

The architecture should specify which systems are trusted immediately, which require checkout, which reserves stay isolated, and which tasks are deliberately deferred. A rapid race to 'activate the base' can destroy diagnostic evidence or consume protected emergency resources.

Mission closure evidence matrix

Evidence matrix	Power	Water	Air	Mobility	Medical
Installed	☐	☐	☐	☐	☐
Available	☐	☐	☐	☐	☐
Verified	☐	☐	☐	☐	☐
N-1	☐	☐	☐	☐	☐
Recoverable	☐	☐	☐	☐	☐

Board gate

GO / HOLD / NO-GO from evidence

Critical functions are checked across installed, available, verified, N-minus-one and recoverable states before a GO/ HOLD/NO-GO board. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this dossier; not a mission-certified drawing.

A 500-sol campaign needs a maintenance economy

Long-duration missions accumulate wear, deferred work, software changes, spares consumption and human knowledge. Maintenance is not a sequence of isolated repairs; it is an economy of limited parts, tools, specialist time and access windows. Every non-critical repair that is deferred creates a backlog that can later combine with an unrelated failure.

The capstone should therefore track corrective maintenance, preventive maintenance, deferred defects and configuration changes. A system may be technically operable while its maintenance debt is rising toward an unacceptable common-cause risk.

Rescue and abort need geometry, time and authority

An abort option that exists only as a sentence is not an option. It needs a vehicle, reachable location, energy or propellant, crew state, route, weather or environment, communications and a decision authority that can act before the option expires.

The board should test rescue routes under degraded assumptions: one rover unavailable, one airlock isolated, one communications path down, a medical casualty slowing movement, or a storm reducing power. The purpose is not to predict every emergency but to discover whether several failures consume the same supposed escape path.

Communication loss tests autonomy rather than patience

Primary-source bridge: NASA — Human Spaceflight.

A 48-hour Earth–Mars communications outage should not reduce the crew to waiting. The mission needs local authority boundaries, cached procedures, engineering data, medical decision support, software rollback capability and criteria for actions that must remain reversible until contact returns.

The handover between local and Earth-based authority must also be explicit. When communications resume, the crew should not need to reconstruct two days of decisions from memory; the event log, configuration state and unresolved risks become part of the recovery evidence.

Return and inheritance are part of mission closure

The return vehicle, ascent path, rendezvous sequence and consumable reserve cannot be treated as a distant final chapter. Their readiness is a protected capability throughout the surface campaign. Maintenance or cannibalisation decisions must not unknowingly consume return-critical parts.

A successful mission also leaves an inheritance package: calibrated equipment, configuration records, known defects, remaining inventory, environmental baselines, software versions and lessons learned. The next crew should inherit evidence, not folklore.

Qualification casebook — six board decisions

1. **1. Cargo landed, but ISRU has only run for six hours.** The plant met nameplate output once; no long-duration quality run exists.

Reasoned disposition — open after making your own decision

HOLD crew dependence on the plant. Count stored verified product, not optimistic future production, until endurance and quality evidence exist.

2. **2. A noncritical pump repair is deferred for the third time.** The system remains redundant today.

Reasoned disposition — open after making your own decision

Record the maintenance debt and the shared specialist/spares dependency. Repeated deferral can convert redundancy into a latent common-cause vulnerability.

3. **3. One rover is unavailable during a distant science campaign.** The remaining rover can physically carry the team.

Reasoned disposition — open after making your own decision

Recompute rescue geometry and protected return margin. Capacity to continue science is not the same as capacity to preserve rescue options.

4. **4. Earth communications disappear for 48 hours.** No life-threatening failure is present.

Reasoned disposition — open after making your own decision

Continue under predefined local authority, preserve reversible choices where possible, log configuration changes and prepare a structured handover for restored contact.

5. **5. A spare intended for the return vehicle can repair the habitat.** The habitat fault is serious but not immediately life-threatening.

Reasoned disposition — open after making your own decision

Use the integrated risk board: consuming return-critical inventory may solve one problem by deleting an abort path. The decision requires explicit reclassification of return capability.

6. **6. The final mission report shows all budgets green.** Several green budgets rely on the same two specialists working overtime.

Reasoned disposition — open after making your own decision

Re-open closure. Shared human workload can invalidate apparently independent margins; the board must prove sustainable staffing and backup competence.

Mastery studio — four extended review problems

Use these problems as an integrated mission board: follow dependencies across phases and show which protected capability is spent by each decision.

1. **1. Precursor acceptance review.** The precursor power system has operated for 200 hours, but all testing occurred under mild seasonal conditions. Crew departure is six weeks away. Decide whether the evidence is sufficient.

Extended reasoned answer — open after attempting the problem

Treat runtime as only one dimension of verification. Compare the tested environmental envelope with the conditions expected before and after crew arrival, including dust, low-temperature periods, storage state and likely load combinations. If critical cold or degraded-power cases have not been exercised, the evidence may support partial availability but not full mission-ready status. The board can close the gap with targeted tests, additional stored energy, alternate power paths or a departure HOLD. The objective is not to demand infinite testing; it is to know which untested condition could defeat a crew-critical function.

2. **2. Maintenance debt under a green dashboard.** Three noncritical defects are open, each individually tolerated by redundancy. They all require the same specialist and share access to one equipment bay. Explain why the mission may still be drifting toward HOLD.

Extended reasoned answer — open after attempting the problem

The defects are not independent once they share human expertise and maintenance access. A fourth failure in the same bay could force several repairs into the same window, while the specialist's available hours become the common bottleneck. Add the open work to the workload and configuration ledger, identify which redundant paths are actually protected, and decide whether any deferred task should be pulled forward before the campaign enters a more demanding phase. A green functional dashboard can hide maintenance debt that is consuming future recovery capacity.

3. **3. Emergency cannibalisation decision.** A habitat cooling fault can be repaired immediately using a valve reserved for the ascent vehicle. No equivalent spare exists on site.

Extended reasoned answer — open after attempting the problem

Do not frame the choice as habitat versus spare. Identify the time to loss of habitat function, alternate cooling modes, whether the ascent valve can be replaced before the return window, and what abort capability is lost if it is consumed. If the habitat threat is immediate and no other recovery exists, cannibalisation may be justified, but the architecture must then explicitly reclassify ascent readiness and create a replacement plan. The decision record should show which protected margin was spent and what new mission state results.

4. **4. Communication blackout and medical uncertainty.** During a 48-hour Earth blackout, a crewmember develops symptoms that are concerning but not immediately life-threatening. The local crew has diagnostic equipment but limited specialist support.

Extended reasoned answer — open after attempting the problem

Operate within pre-authorised local medical and command boundaries. Use available diagnostic protocols, trend vital evidence, preserve reversible choices and document every intervention. The commander and medical officer should know which thresholds trigger evacuation, isolation or use of scarce medication without waiting for Earth. When contact returns, transmit a structured timeline, measurements, treatments and unresolved questions rather than a narrative reconstructed from memory. The exercise tests whether the mission has real autonomy, not whether the crew can simply wait.


Primary sources used in this qualification dossier

- NASA — Moon to Mars Architecture
- NASA — Environmental Control and Life Support System
- NASA — Human Spaceflight

Closure standard. The learner can defend an integrated mission architecture across verification, workload, rescue, autonomy and return without hiding shared dependencies behind green subsystem slides.

Mission assurance rehearsal — make cross-system decisions across four time horizons

The capstone now contains budgets, failure injections and GO/HOLD/NO-GO logic. The operational layer adds the layer that senior operators actually need: one failure can be survivable for eight hours but mission-threatening after thirty days, and a locally attractive repair can destroy the return option. The learner therefore practices cross-system decisions on linked time horizons rather than closing each subsystem in isolation.



Decision atlas — Mission assurance rehearsal — make cross-system decisions across four time horizons. Pedagogical synthesis by Delta-Sierra; use the full-size link for fine labels.

Separate survival, stabilization, recovery and return horizons

The first horizon is immediate survival: breathable atmosphere, pressure, fire control, thermal protection, water and medical stabilization. The second horizon is stabilization over roughly hours to days: diagnose the fault, protect buffers, control crew workload and establish a repair path. The third horizon is sustained recovery: restore redundancy, rebuild inventory, retire maintenance debt and prove stable operation. The fourth horizon preserves the return mission and the next crew. A decision that is correct in the first horizon can be unacceptable if it irreversibly damages a later one.

This structure forces the review board to state what ‘success’ means at each time. Running a spare pump may restore water today but consume the only unit reserved for the ascent vehicle. Cannibalizing a rover may protect habitat power yet eliminate medical rescue radius. The architecture needs explicit authority for such trades before the emergency, not improvised moral debate after alarms start.

Cross-system dependencies deserve their own ledger

Power, thermal rejection, communications, life support, mobility, medical capability and crew time form a network. Redundancy in one box does not protect the mission if two nominally independent branches depend on the same coolant loop, software service, electrical bus or specialist. The mission assurance ledger therefore records not only components but shared dependencies and recovery resources.

A useful review question is: which single failure would make three teams simultaneously request the same person, power feed or spare? That question exposes common causes that are invisible in a subsystem block diagram. NASA Moon to Mars Architecture is used as a primary source bridge for system-of-systems thinking; the specific Delta-Sierra ledgers remain pedagogical review tools.

Crew workload must be budgeted during degraded operation, not only nominal duty

A failure often increases manual monitoring, sampling, maintenance, medical observation and reporting at the same time that sleep and cognitive performance are degrading. The mission therefore needs a degraded-mode workload budget: critical tasks per shift, required qualifications, backup personnel and tasks that can be deferred. If one specialist exceeds the sustainable workload, that human becomes a common-cause failure point.

The workload board should include handover time and the cost of context switching. Ten five-minute checks distributed through a shift may be more disruptive than one fifty-minute task because they fragment sleep, maintenance and planning. A resilient architecture buys automation where automation reduces sustained workload, but it also preserves manual control when automation itself is the fault source.

Medical capability is a function, not a cabinet of supplies

A medicine list does not prove medical capability. The architecture must connect likely conditions to diagnostic tools, medications, procedures, isolation space, communications support and trained people. A stock may exist yet be unusable because the crew lacks the skill, power, sterility, imaging or monitoring needed to apply it safely.

The capstone therefore asks whether medical capability remains coherent during the same failures that affect power, water and communications. NASA ECLSS material provides a primary bridge for the environmental functions that medical survival depends on; crew health decisions still require dedicated medical standards and professional protocols beyond this course.

Communication loss is a test of delegated authority

A Mars crew cannot treat Earth as a real-time control room. Communication delay already separates advice from action, and an outage can extend that separation. The mission must predefine which decisions the crew can take autonomously, what evidence must be recorded, which actions are reversible and which commitments require waiting when waiting is safe.

This is not a political question disguised as engineering. Authority affects response time. If a pressure leak requires immediate isolation, the crew must not need a terrestrial vote. If the proposed repair would consume the last return-critical component and the habitat remains safe for forty-eight hours, deliberate consultation may be appropriate. The decision matrix therefore couples urgency, reversibility and strategic consequence.

Spares and cannibalization create an option portfolio

A spare is valuable because it preserves a future repair. Once installed, that option disappears. Cannibalizing another system can create a repair today by deliberately opening a risk tomorrow. The mission should therefore track not only spare count but which functions each spare can protect and which future failures become unrecoverable after use.

An option portfolio can rank parts by cross-system criticality, lead time, repairability and substitution. The board does not need a mathematically perfect economic model; it needs visibility. If a component protects both water processing and return-vehicle thermal control, using it for a noncritical convenience fault should require explicit review.

Maintenance debt is a mission state variable

Deferred inspections, temporary bypasses, quarantined components and partially restored redundancy accumulate as maintenance debt. The mission may still appear green because every critical function is currently operating, but the architecture has less ability to absorb the next failure. A board that reports only present availability hides this loss of resilience.

The course therefore treats maintenance debt as an explicit condition for campaign pace. Science sorties, construction and ISRU expansion can be slowed until deferred work falls below a defined threshold. This is a strategic choice: accepting lower short-term productivity to restore the margins that protect the remaining months of the mission.

The return vehicle must remain outside casual trade space

The return system carries propulsion, power, thermal, life-support and guidance resources that may look attractive during a surface emergency. The architecture should predefine which of those resources are truly shareable and which are protected because consuming them can strand the crew. Emergency authority may override that protection to save lives, but the consequence must be explicit.

This keeps the capstone honest. 'We can always use the return vehicle' is not a free contingency; it is a conversion of one survival path into another. The board must then rebuild or replace the lost return capability before declaring mission recovery. NASA human spaceflight material is used as a primary bridge for integrated crewed-mission constraints; the scenario values in this course remain pedagogical.

Handover closes the mission across generations of crews

A settlement mission is not complete when the first crew survives. The next team inherits equipment age, software versions, maintenance debt, local climatology, medical history, resource inventories, known weak interfaces and unresolved anomalies. A good handover turns experience into institutional memory instead of forcing each crew to rediscover risk.

The final capstone dossier therefore includes a configuration baseline, open-risk register, spares status, trend plots, consumable forecasts, known procedural deviations, training gaps and the evidence behind every current operating limit. This is the difference between a heroic expedition and a durable operating system.

Operational review board — six decisions to defend

1. **1. Water loop stable for eight hours, no redundant pump.** The habitat can meet current demand but another pump fault would end potable-water recovery.

Reasoned disposition — open after making your own decision

Classify survival as green, recovery as HOLD. Reduce discretionary demand, accelerate repair/redundancy restoration and do not expand operations until a second credible path exists.

2. **2. Rover cannibalization would restore habitat cooling.** The rover is also the medical rescue asset.

Reasoned disposition — open after making your own decision

Compare habitat thermal time-to-limit with rescue dependence and alternative cooling actions. If cannibalization is necessary, formally declare the lost rescue option and create a replacement plan before routine EVA resumes.

3. **3. One operator owns every critical workaround.** Manual monitoring keeps three systems stable.

Reasoned disposition — open after making your own decision

Treat human workload as common cause. Cross-train, automate or reduce mission tempo before fatigue turns a recoverable technical fault into an operational cascade.

4. **4. Medical imaging fails during a communications outage.** The crew has a stable but uncertain abdominal case.

Reasoned disposition — open after making your own decision

Use pre-authorized medical protocols, protect communications attempts and evaluate transfer/observation options. Do not pretend that a well-stocked pharmacy replaces lost diagnostic capability.

5. **5. A return-critical spare could restore nonessential science power.** Science loss is politically painful but not life-threatening.

Reasoned disposition — open after making your own decision

Protect the return-critical option. The mission can accept science downtime more safely than it can accept an unrepairable return chain.

6. **6. Maintenance debt crosses the board threshold.** Several temporary bypasses remain while the schedule calls for a major construction push.

Reasoned disposition — open after making your own decision

Slow the campaign and retire debt before adding new failure exposure. Productivity is not a margin if it consumes the system's remaining resilience.

Mission rehearsal notebook — reason through evidence before revealing the disposition

Integrated exercise A — 18-hour power deficit during medical isolation

A crew member is isolated for a suspected infectious condition while a dust event cuts solar generation and one battery string is unavailable. The mission board must avoid solving each problem independently. Medical isolation may increase ventilation, monitoring and crew-time demand exactly when power and thermal margins are shrinking. The board first establishes the immediate survival load: pressure control, air revitalization, critical medical equipment, communications and minimum thermal control. It then identifies loads that can be deferred, the battery endurance under the reduced configuration and the expected recovery of generation. The medical team states which diagnostic or treatment functions are time-critical and which can wait. Operations protects sleep and handover because exhausted crews are themselves a common cause. The final decision may be to reduce industrial loads, suspend EVAs, place noncritical habitat zones in a low-power state and preserve a defined reserve for a second failure. The exercise demonstrates why integrated margins are more useful than a collection of subsystem green lights.

Integrated exercise B — water processor fault with a viable but labour-intensive workaround

The primary water processor loses a controller. A manual bypass can maintain safe water quality, but it requires sampling every two hours and filter intervention twice per shift. At first glance the workaround closes the water problem. The capstone board must go further: which crew qualifications are required, how does the new workload interact with maintenance, medical duties and sleep, and how long can the bypass be sustained before human performance becomes the limiting resource? The team uses verified tank inventory to buy diagnostic time, cross-trains an additional operator and schedules a controller repair. It establishes a deadline before the workaround becomes normalised. Recovery is not declared when water flows; it is declared when qualified output, sustainable workload and redundancy are restored. This is a classic example of technical recovery creating operational debt that must be retired before mission tempo returns to normal.

Integrated exercise C — rover loss removes the surface rescue option

A mobility failure strands the long-range rover in a repair state near the habitat. No crew is endangered immediately, and the habitat itself is healthy. The hidden consequence is that the mission has lost part of its rescue architecture for future EVAs. The board therefore does not classify the event as merely a transportation inconvenience. It freezes activities whose emergency return depended on that rover, checks alternate vehicles and walking/refuge geometry, identifies spares and repair time, and updates the risk register. If a construction objective depends on reaching a distant site, the schedule slips even though all life-support systems are green. This teaches the learner to evaluate functions rather than hardware labels: a rover can be simultaneously transportation, logistics, medical evacuation and contingency power. Restoring one driving mode may not restore all those functions.

Integrated exercise D — return propellant reserve competes with surface survival

A prolonged habitat fault creates a proposal to use a return-stage resource that was intended to remain protected until departure. The board has to separate moral urgency from technical ambiguity. If using the resource is necessary to prevent imminent loss of life, the decision can be obvious, but it must be recorded as a deliberate conversion of the return architecture. If the habitat can remain safe while another repair is attempted, the board protects the return reserve and uses the available time. In either case, recovery cannot later be declared merely because the habitat is stable. The team must rebuild a credible return path, revise mission duration if necessary and update the evidence supporting departure readiness. This exercise is designed to prevent a common narrative shortcut: treating a return vehicle as an infinite contingency store. Every shared resource has an opportunity cost, and the capstone must show that cost explicitly.

Primary-source bridges used in this operational rehearsal

- NASA — Moon to Mars Architecture
- NASA — Environmental Control and Life Support Systems
- NASA — Humans in Space

Closure standard. The learner can explain the evidence chain, identify the limiting uncertainty or dependency, defend a degraded-state decision and state what a real mission would still need to verify.

Mission integration control room — protect options across a 500-sol campaign

A mission architecture is not closed merely because each subsystem has a plausible design. It is closed when the team can explain how the functions interact under time pressure, which reserves are protected, which degraded states are acceptable, who may spend a protected margin, how the return path is preserved and what evidence is required to move from survival to stabilization, then recovery and finally full mission capability.

This control-room dossier is deliberately cross-system. It treats power, water, atmosphere, thermal control, mobility, medicine, crew time, spares, communications and return readiness as a coupled network. The central habit is to ask which function will be lost first if no action is taken. That time-to-loss-of-function can be minutes for atmosphere, hours for a rescue opportunity, days for a water reserve or months for some maintenance debt. The loudest alarm is not automatically the shortest clock.

Mission integration board — shared dependencies change what is actually urgent

P = primary dependency · S = supporting dependency · ! = immediate common-cause concern

protected function	power	thermal	water	crew time	spares	time to loss of function
breathable atmosphere	P	S	S	S	S	minutes–hours safe-haven clock dominates
qualified water	P	S	P	S	S	hours–days inventory buys diagnosis time
medical capability	S	S	S	P	P	condition-dependent clinical clock can be shortest
surface rescue	P	S	S	P	P	minutes–hours distance and weather matter
return readiness	P	S	S	P	P	days–months unless a protected margin is spent

The board acts on the shortest credible loss-of-function clock, not on the subsystem with the loudest alarm.

Delta-Sierra pedagogical synthesis · actual mission clocks require verified inventories, crew state and configuration data.

Mission-integration atlas — shared dependencies and the shortest credible loss-of-function clock determine what is actually urgent.

1. Define protected functions before discussing hardware

Hardware names are useful for configuration control, but mission decisions protect functions. “Water processor A is down” is a hardware statement. “Qualified drinking-water production is unavailable, 180 litres remain accessible, manual transfer is available and the present demand empties the protected reserve in 38 hours” is a mission statement. The second form shows time, evidence and options.

The control room therefore maintains a list of protected functions: breathable atmosphere, safe thermal environment, qualified water, medical capability, safe-haven capacity, surface rescue, communications/ autonomy, food security, critical maintenance, and return readiness. A single component can support several functions. A rover can be mobility, rescue and contingency power. A power bus can simultaneously support air revitalization, thermal control and communications. This is why a component-by-component risk list can hide common-cause exposure.

2. Time-to-loss-of-function is the clock that orders decisions

When multiple faults occur, the first question is not “which subsystem has the highest criticality label?” It is “what is the earliest credible loss of a protected function if we do nothing?” The answer should include the assumptions that create the clock. A life-support buffer may buy eight hours only if demand remains at the present level. A medical condition may shorten the rescue clock. A dust event may extend rover travel time and therefore change the time available to recover a stranded crew.

The clock also changes when the team acts. Reducing non-essential demand can extend water or power endurance. Moving crew to a smaller safe-haven volume can change atmospheric and thermal loads. Cancelling an EVA can protect both medical capacity and mobility. These are not free resources; every action has consequences elsewhere. The board records the new clock after each major action rather than continuing to quote the original estimate.

3. Mission inventory has states: onboard, accessible, qualified, available and protected

A tonne of water somewhere in the architecture is not automatically a tonne the crew can drink. It may be in a tank that cannot be transferred, in quarantine pending quality results, reserved for another function, frozen in a process line, or physically available but protected for return or emergency use. The same logic applies to oxygen, batteries, spares, pharmaceuticals and propellant.

For every mission-critical inventory, the review board should be able to distinguish: total physically present; accessible with the current configuration; qualified for the intended use; available to spend without violating another commitment; and protected by policy. This vocabulary prevents double-counting. It also makes a controversial decision visible: spending a protected reserve is a deliberate change to the architecture, not a bookkeeping trick.

Primary bridge: NASA — Environmental Control and Life Support Systems. The operational inventory-state vocabulary in this dossier is a Delta-Sierra teaching construct, not an official NASA taxonomy.

4. Common-cause dependencies should be drawn before the emergency

Nominal redundancy is weak if two “independent” systems share the same upstream dependency. Two water-processing trains can share power, cooling, software, calibration standards or the same two technicians. Two rovers can share a charging station. Two communications paths can share an antenna pointing service or software build. A settlement can therefore have duplicated boxes while retaining one hidden common cause.

The control-room ledger marks the shared dependencies explicitly and asks which one could defeat more than one protected function. The purpose is not to make the diagram pessimistic; it is to identify where diversity, isolation, manual workarounds or protected spares actually buy resilience. When a common cause is discovered during a failure, the lesson must be fed back into the architecture rather than treated as an unusual one-off event.

5. Workload is a consumable during degraded operation

A workaround that requires two specialists to perform manual sampling every hour may technically restore a function while simultaneously exhausting the crew. The mission must therefore budget degraded workload just as it budgets water or energy. The ledger records required qualification, people per task, minutes per shift, duration, handover burden and which other work must stop.

Workload also interacts with error probability. A plan that depends on repeated manual valve alignment, calculation and sampling during sleep loss can create new failure paths. The board should deliberately simplify the configuration where possible and assign independent verification to the most hazardous manual steps. A workaround that is sustainable for four hours may be unacceptable for four days.

6. Maintenance debt can preserve today while destroying next month

Deferred maintenance is sometimes a rational emergency decision. If a processor must be kept running to restore a reserve, a non-critical inspection can be postponed. The error is to let the deferred work disappear from the mission state after the immediate crisis. The framework therefore treats maintenance debt as a ledger with due task, risk if deferred, required skill/spare, earliest safe opportunity and interaction with return readiness.

Debt can become common cause. A crew may repeatedly cannibalize the same class of spare, defer software qualification and use temporary jumpers until the settlement contains many individually understandable exceptions. The combined configuration then becomes difficult to reason about. Recovery therefore includes configuration simplification and debt reduction, not only restoration of nominal output.

7. Medical capability is a chain of functions under a clock

Remote medicine on Mars cannot be modelled as a cupboard of supplies. It requires recognition, monitoring, decision support, procedures, trained crew, sterile capability where needed, power, communications when available, pharmaceuticals and the ability to keep the rest of the habitat safe while care is delivered. A medical event can also consume mobility, crew time and safe-haven capacity.

The integration board does not make clinical decisions from an engineering table. It ensures that the clinical team can state what capability is currently available and which system faults threaten it. If a cooling fault risks both medical storage and avionics, the shared dependency belongs in the mission ledger. If a patient cannot tolerate a delayed rover rescue, the rescue clock becomes the governing operational constraint.


Primary bridge: NASA — Humans in Space. Mission-specific medical standards, crew selection and treatment protocols remain outside the authority of this pedagogical scenario.

8. Communications loss tests delegated authority, not patience

A Mars crew cannot treat Earth as a real-time controller. The architecture should define decisions the crew can make locally, decisions that require later notification, and decisions that should wait when waiting is safe. The authority model must also survive ambiguity: if communications fail during a surface emergency, the crew needs predefined objectives and protected constraints, not a list of commands that assumes Earth can respond immediately.

Delegated authority is strongest when it is bounded by evidence. “Do whatever is necessary” is not a useful rule. “Protect life, preserve the return branch unless immediate survival requires its use, maintain a written assumption ledger, and record any protected margin spent” gives autonomy while preserving reviewability. Earth can then reconstruct why a decision was made after communications return.

9. Return readiness is a parallel branch throughout surface operations



Mission-integration atlas — the return path remains a tracked mission function while surface operations consume resources and accumulate maintenance debt.

The return vehicle and its supporting resources should not become a casual source of surface spares, power or consumables merely because they are nearby. Some cross-use may be intentionally designed, but the decision must identify what return capability is lost and how it will be restored. The return branch includes more than propellant: vehicle configuration, power/thermal health, navigation, critical spares, crew medical fitness, departure-site access, communications and the procedures needed to execute the departure.

If survival requires consuming a protected return resource, the correct response may absolutely be to use it. The discipline is to declare that the mission state has changed. The team then rebuilds a credible return architecture or accepts that the original return objective no longer exists. This prevents optimistic accounting in which the same resource is simultaneously spent on the emergency and still counted in the departure reserve.

10. Integrated failure campaign — nine injects for one control room

Inject 1 — power deficit and rising CO₂ occur during a dust event

Reasoned disposition. Rank the atmospheric loss-of-function clock against battery endurance and thermal constraints. Shed non-essential electrical loads, protect air revitalization and monitoring, move crew if a smaller controlled volume gives more margin, and record the new battery/atmosphere clocks. Do not optimize solar recovery assumptions while the immediate breathable-atmosphere clock is shorter.

Inject 2 — qualified water reserve is healthy, but the only specialist who can maintain the manual workaround is injured

Reasoned disposition. Recalculate endurance using the degraded staffing configuration. Train/brief a backup if time permits, reduce demand, and decide whether the system should be simplified rather than maintained through a fragile specialist-dependent workaround. Crew qualification is a mission resource.

Inject 3 — a rover is serviceable for logistics but not certified for emergency rescue

Reasoned disposition. Keep the two capability states separate. A vehicle that can move cargo may lack life-support support, communications, range margin or medical accommodation needed for rescue. Restore or prove the rescue function before counting it in the contingency architecture.

Inject 4 — two nominally redundant processors fail after the same software maintenance event

Reasoned disposition. Treat the shared software/configuration path as common cause. Isolate configurations, roll back only with a verified package, protect buffers and obtain independent output-quality evidence before reconnecting the loop. Hardware duplication did not provide functional independence.

Inject 5 — a medical emergency consumes the only pressurized rover during a maintenance campaign

Reasoned disposition. Medical rescue may dominate, but the board must immediately recompute the other functions that depended on that rover: remote rescue, cargo, spare delivery and perhaps contingency power. Suspend operations whose safe-return geometry depended on the removed vehicle.

Inject 6 — Earth communications are lost while a protected reserve may need to be spent

Reasoned disposition. Use the delegated-authority rules. If survival requires the reserve, spend it and document the changed return architecture. If the crew remains safe with alternatives, preserve the reserve until the evidence justifies conversion. Lack of Earth contact does not erase local decision criteria.

Inject 7 — water, food and power margins are individually positive but all depend on the same maintenance specialist

Reasoned disposition. The shared human dependency converts three apparently separate margins into one common cause. Reallocate training, simplify maintenance, protect rest and prioritize the specialist's tasks by time-to-loss-of-function. The architecture is not robust until competence is distributed or the dependency is reduced.

Inject 8 — a temporary repair restores output but creates untracked configuration exceptions

Reasoned disposition. Record the temporary configuration, inspection interval, limits and retirement plan. Function restored is not mission restored. The repair enters maintenance debt and must be cleared or formally accepted before it can become invisible background configuration.

Inject 9 — the habitat is stable, but return vehicle maintenance and crew conditioning have slipped for six weeks

Reasoned disposition. Do not declare full recovery from surface stability alone. Return readiness is a protected parallel function. Rebuild the vehicle/crew departure state, identify the next departure opportunity and decide which surface activities must pause until the return branch is back inside its acceptance envelope.

11. The integration evidence pack should survive a crew change

Artifact	Minimum content	Failure it prevents
Assumption register	Value, source, uncertainty, owner, expiration/review trigger	Hidden premises becoming permanent facts.
Protected-margin ledger	Inventory, accessible/qualified amount, policy, authority to spend	Double-counting reserves.
Dependency map	Shared power, cooling, software, sensors, crew and spares	False redundancy.
Time-to-loss ledger	Function, present clock, assumption, action that changes it	Reacting to the loudest alarm rather than the shortest clock.
Maintenance-debt register	Deferred work, risk, due point, temporary configuration	Emergency workarounds becoming undocumented normality.
Return-readiness board	Vehicle, consumables, navigation, crew, site access, critical spares	Surface success consuming the return path.
Decision log	Evidence known, alternatives, authority, protected margin spent, review trigger	Later teams being unable to reconstruct why a risk was accepted.

12. Logistics closure means proving access, not merely counting mass

A spare on Mars can still be unavailable to the function that needs it. It may be stored in an unpressurized cache during a dust event, behind an inoperable hatch, compatible with a different revision, or require a tool and technician who are both committed elsewhere. The logistics ledger therefore records location, compatibility, access conditions, preservation state, required tools and the time needed to deliver the item to the failure site.

This turns “we have three spares” into a defensible statement. The board can distinguish physical stock from usable stock and can detect when multiple functions depend on the same last spare. It also reveals when local manufacturing can substitute for a low-complexity part and when it cannot replace a certified component whose material, tolerances or test history are essential.

13. Food and water security should be expressed as recoverable mission functions

Food mass is necessary but not sufficient. The mission needs nutritional adequacy, storage integrity, preparation capability, crew acceptance and a plan for losses. Locally grown food can reduce dependence but adds lighting, water, nutrient, labour and contamination dependencies. It should therefore be credited by demonstrated delivery, not optimistic crop area.

Water is similarly multi-state. Potable reserve, hygiene water, process water and water trapped in another subsystem do not have identical mission value. In an emergency, some streams may be reallocated, but quality and downstream consequences must be explicit. The mission integration board should show how long essential drinking/medical demand can be met before optional uses are considered.

14. Power restoration needs a black-start story

A settlement can possess enough installed generation and storage yet be unable to restart after a major bus collapse. Some loads are needed to bring other generators, thermal loops, controls or communications back online. The architecture therefore identifies black-start-capable sources, the order of restoration, inrush constraints and the minimum instrumentation needed when the main network is unavailable.

The integration lesson is cross-system: power restoration may require thermal control, and thermal control may require power. A restart sequence breaks that circular dependency with protected low-power modes, local controllers, manual valves or isolated microgrids. The evidence pack should show that this sequence was tested in degraded conditions rather than inferred from nominal schematics.

15. Surface mobility is part of the emergency architecture

Rovers support science and logistics, but the capstone must separately prove rescue range, route accessibility, communications, life-support support and degraded travel time. A vehicle that can drive 30 kilometres nominally may not provide 30 kilometres of rescue radius when the return energy reserve, terrain, dust, crew condition and possible second fault are included.

The control room therefore protects at least one credible rescue path whenever crews are outside immediate habitat reach. If two teams use the same rover reserve or the same charging point, their missions are coupled. The schedule should reveal that coupling before departure rather than discovering it after two simultaneous problems.

16. Crew competence is an architecture with depth and expiration

A long mission cannot assign one irreplaceable expert to every critical function. Skills decay, illness happens and simultaneous failures compete for the same person. The training matrix therefore records primary operator, backup, last practice, required supervision and which procedures can be executed from checklists by a non-specialist. Some tasks may require two qualified people because independent verification is part of the safety case.

The objective is not that every crew member becomes equally expert. It is that the mission can maintain essential functions through foreseeable absences and workload peaks. Training time is itself a resource, so the architecture chooses which competencies need deep redundancy and which can rely on slower remote support when communications permit.

17. Consumable policy should distinguish routine use, contingency use and protected reserve

If every team can draw from the same generic reserve, the mission may discover that several independent plans all counted the same resource. The framework therefore recommends policy labels: routine allocation; contingency allocation requiring a local operations decision; protected reserve requiring mission-level authority; and inaccessible/committed stock. The labels are pedagogical, but the accounting principle is fundamental.

A protected reserve can still be spent. The purpose of protection is to force the decision to show what future function is being traded away. That record becomes part of the revised mission architecture. It also prevents later reports from accidentally presenting the pre-emergency reserve as if it still existed.

18. Long-duration human performance belongs in the mission schedule

Fatigue, sleep disruption, workload, interpersonal strain and reduced attention can alter the reliability of every manual procedure. The architecture therefore needs protected rest, realistic shift handovers, limits on consecutive high-risk tasks and a way to reduce mission tempo during prolonged degraded operations. These are not optional wellness benefits; they influence the probability that complex maintenance and medical actions are performed correctly.

A sustained emergency can be more dangerous after the first successful repair because the crew is tired while the system remains fragile. The recovery gate should therefore include human state: enough rested qualified people must exist to operate the repaired configuration and respond to the next fault.

19. Science and exploration must be schedulable around maintenance reality

A mission that allocates all nominal crew time to science and exploration has already hidden maintenance. Preventive work, inspections, training, housekeeping, software/configuration management, inventory, medical checks and anomaly follow-up consume real hours. The capstone schedule therefore includes these categories before it claims spare capacity for discretionary activities.

When maintenance debt rises, science is not automatically cancelled; the board weighs value, risk and timing. But the trade must be visible. A high-value time-sensitive observation may justify carrying a small amount of debt, while a routine traverse may be delayed because it would consume the only rescue rover during a fragile recovery period.

20. Six additional integration exercises

Exercise A — a spare exists, but it is stored outside during a dust event

Count the spare as physically present but not immediately accessible. Estimate the safe retrieval window, alternate repair, and time-to-loss-of-function. If the function clock is shorter than safe access, the architecture needs another option; inventory mass alone does not solve the fault.

Exercise B — greenhouse output is nominal but crop disease forces quarantine

Remove quarantined food from qualified inventory and recalculate nutritional endurance. Protect seed/clean-zone capability, investigate the contamination path and decide how much lighting/water should continue to be spent on the affected crop. Agricultural production is not credited until it is usable.

Exercise C — the main power bus collapses after sunset

Use the black-start plan: protected source, minimum controls, communications, thermal/life-support loads, then staged restoration. Do not reconnect all loads because total stored energy appears sufficient. Inrush and coupled restart dependencies can cause a second collapse.

Exercise D — two critical maintenance tasks require the same expert during a medical emergency

Order the tasks by time-to-loss-of-function, use qualified backups for bounded subtasks, suspend lower-value operations and protect rest. If neither task can be safely staffed, the mission architecture has insufficient skill redundancy and must enter a more conservative operating state.

Exercise E — using a return-system battery would solve a surface power shortage

Identify the immediate threat, alternatives, amount of protected margin spent and how return readiness changes. If the surface threat is life-critical, spending the battery may be correct. But the return branch must then be re-qualified or the mission must accept a changed departure capability.

Exercise F — science schedule is on target while maintenance debt doubles

Make debt visible in the mission board. Identify which deferred tasks threaten common dependencies or return readiness, then rebalance the schedule. Meeting the science plan cannot compensate for silently degrading the infrastructure that keeps the mission viable.

21. Final acceptance matrix — what “mission closed” should mean

Domain	Evidence before approval	Unacceptable shortcut
Life support	Verified production, qualified reserves, degraded/manual mode, restart proof	Nominal recovery percentage alone
Power/thermal	N–1 or accepted degraded state, black-start sequence, critical-load priority	Installed capacity exceeds average load
Mobility/rescue	Range with reserve, crew support, route, charging and alternate plan	Rover can drive the nominal distance
Medical	Capability chain, trained crew, supplies, environment, decision authority	Inventory of medicines
Maintenance	Backlog, skills, spares, temporary configurations, debt retirement	All major boxes currently running
Autonomy	Delegated authority, offline procedures, evidence/decision log	Earth can always advise
Return	Vehicle, crew, navigation, protected resources, access and next opportunity	Return assets still exist physically
Handover	Another crew can reconstruct state, margins, debts and open risks	Original team remembers the rationale

Primary bridges for this integrated dossier: NASA Moon to Mars Architecture, NASA ECLSS, and NASA Humans in Space. The ledgers and gates above are pedagogical synthesis tools, not an assertion that NASA uses these exact Delta-Sierra labels.

Mission integration mastery board — prove that the whole architecture can survive the campaign

A final mission architecture is not validated by adding up the best-case performance of individual subsystems. The integration board must show that protected functions survive the actual sequence of mission phases, that failures do not consume the same hidden reserves twice, that people can execute the degraded procedures, and that the return branch remains viable after local recovery actions. This section therefore treats the settlement and vehicle as one coupled system of systems whose technical state, crew state, inventory, maintenance condition and decision authority evolve together.

Competency 1 — define protected functions before choosing hardware

The architecture should begin with functions that must remain available: breathable atmosphere, potable water, thermal control, critical electrical power, medical capability, shelter, communications, mobility/rescue and the ability to preserve or restore the return path. Hardware is then selected to deliver those functions with measurable margins. This ordering matters because two pieces of equipment that look redundant can share a power bus, coolant loop, software image, consumable or specialist and therefore fail together.

For each protected function, the board should know what “available” means. Installed capacity is not the same as currently usable capacity; usable capacity is not the same as verified capacity after maintenance; verified nominal capacity is not the same as N–1 capability after a credible failure. These states should be visible in the mission ledger so a subsystem cannot be declared green merely because hardware exists on site.

Competency 2 — map dependencies until common causes become visible

A dependency map links every protected function to the resources, interfaces and people it requires. Water recovery may need power, pumps, filters, sensors, control software and crew maintenance. Power may depend on dust cleaning, thermal management, switchgear and stored energy. Rescue mobility may depend on the same battery inventory and technicians already committed to habitat recovery. A common cause appears whenever apparently independent functions depend on one shared element.

The board should therefore ask a second question for every redundancy claim: “what do both paths still share?” Shared environment can also be a common cause. A dust event can reduce generation, contaminate seals and increase cleaning workload at the same time. A habitat leak can increase gas demand while forcing crew effort into repair and restricting access to adjacent equipment. The architecture must be tested against these coupled events rather than only against neat single failures.

Competency 3 — budget human time as a consumable with qualification constraints

Crew hours are not interchangeable. A task may require a particular medical, electrical, EVA or software qualification. During a failure, monitoring, manual control, troubleshooting and recovery can all demand the same specialists. The board should therefore maintain a workload ledger by role and shift, including backup depth, mandatory rest, handover time and the possibility that an injured or isolated crew member removes a qualification from the available pool.

This is where technical redundancy can fail socially. A backup machine is not a backup function if nobody available can configure or repair it. A procedure that takes four hours in a calm simulation may take much longer in suits, under alarms, after sleep disruption or while another system is unstable. Human-factors margin must therefore be treated as part of mission capacity, not as a soft consideration added after engineering is complete.

Competency 4 — distinguish emergency survival from sustainable recovery

A safe haven, bottled oxygen or battery reserve can keep the crew alive while a system is down, but that does not mean the mission has recovered. Survival is the first horizon. Stabilisation follows when the immediate loss is stopped and the situation no longer deteriorates rapidly. Functional recovery follows when the protected service is restored. Redundancy recovery follows only when backup capability is again available. Mission recovery is later still: maintenance debt, inventory depletion, crew fatigue and return readiness must be restored to acceptable levels.

These horizons prevent premature return to nominal operations. A water loop running through a temporary bypass may deliver potable water but lack redundancy. A power system may meet today's critical load while batteries remain partially depleted and one inverter is unavailable. The correct status is therefore not a binary "fixed/not fixed"; it records which horizon has been reached and which protections remain open.

Competency 5 — protect inventory that is physically accessible, qualified and in the right place

Mass on Mars is not automatically usable inventory. A spare may be in an unpressurised cache inaccessible during the current failure, packaged behind cargo that cannot be moved, or incompatible with the installed configuration. A chemical may exist but not have verified purity for the intended use. A medical stock may be present but expired, temperature-exposed or lacking the diagnostic capability needed to use it safely. Inventory accounting must therefore include location, access time, qualification state and configuration compatibility.

This becomes a design issue for pre-positioned cargo. Critical spares, emergency consumables and recovery tools should be distributed so one fire, depressurisation or blocked corridor does not isolate all of them. The board should deliberately test whether the crew can reach the necessary item under the failure conditions that make it necessary.

Competency 6 — make maintenance debt visible after every contingency

Contingency recovery often borrows resources from the future. A spare is consumed, a redundant channel is cannibalised, preventive maintenance is deferred, a temporary hose remains installed or a crew accumulates fatigue. If those debts are not recorded, the architecture can appear fully recovered while its future resilience has quietly fallen. A maintenance-debt ledger records what was borrowed, the consequence of leaving it unresolved, the resources needed to close it and the latest acceptable closure time.

The same logic applies to software and configuration. A manual override or emergency parameter set may be necessary during recovery, but the board must know whether the system has returned to a controlled configuration. Temporary changes should either be formally accepted into the baseline or removed after their purpose is complete. Otherwise the next crew inherits a system that is physically functional but poorly understood.

Competency 7 — preserve the medical function as a system, not a cabinet

Medical capability depends on more than pharmaceuticals. It includes diagnostic tools, sterile or controlled space, power, water, oxygen, communications or reference material, trained crew, monitoring capability and the ability to isolate infectious or contaminated patients when necessary. A failure in another subsystem can therefore reduce medical capability even when the medication inventory is unchanged.

The mission board should also distinguish stabilisation capability from definitive treatment. A settlement may be able to stabilise a trauma, infection or decompression injury without being able to provide Earth-equivalent definitive care. This distinction changes evacuation, return and risk-acceptance decisions. The architecture should be explicit about which conditions it can diagnose, stabilise, treat and monitor autonomously during communication delays.

Competency 8 — test black-start and cold-start paths, not only steady-state capacity

A system that can carry the load once running may still fail to recover after a shutdown. Pumps, heaters, controllers, valves and network switches may have startup peaks or sequencing constraints. After a major power event, the team should know which loads must come first, which can wait, what communications and instrumentation remain available during the restart and what happens if one startup attempt fails. Black-start is therefore a procedural and configuration problem as much as an energy problem.

This is particularly important for coupled loops. Restarting thermal control may be required before high-power equipment can return; restarting water processing may depend on stable power and sensors; communications may be needed to coordinate remote equipment. A rehearsed sequence with observable gates is more credible than a statement that enough total power exists.

Competency 9 — reserve autonomy for communications loss

Mars crews cannot depend on immediate Earth direction. The mission architecture should define what decisions can be taken locally, what thresholds require a HOLD, what evidence must be logged for later review and which actions are prohibited without broader authority. This is not a philosophical question; it is configuration control under delay. The crew needs enough authority to protect life and mission assets without improvising away safety barriers.

A communications outage should therefore be rehearsed together with another failure, not as an isolated radio problem. If a water leak, medical event or power degradation occurs while Earth contact is unavailable, the crew should already know the decision hierarchy, protected reserves and reporting package that will be sent when the link returns.

Competency 10 — make return readiness a continuously protected branch

Return capability can be eroded gradually by consuming propellant, batteries, spares, crew health margin or vehicle time to solve local problems. The integration board should maintain a separate return-readiness state that is not allowed to disappear inside local subsystem metrics. A local recovery that saves a surface asset but makes the departure vehicle unavailable may be unacceptable unless the mission explicitly decides to cross that threshold.

Return readiness also includes timing. Departure windows, vehicle maintenance, crew conditioning, consumables and surface-to-vehicle logistics have to converge at the right time. The board should therefore ask not only “can the vehicle leave?” but “can the crew and system be brought into departure configuration before the opportunity closes?”

Competency 11 — exercise the architecture with coupled failures and incomplete information

Single-failure tests are useful for verifying components, but integration confidence comes from scenarios where failures interact and information is imperfect. The board should inject combinations that stress shared dependencies: reduced solar generation plus dust-contaminated radiators; water-loop failure plus crew

illness; rover unavailability plus remote maintenance; communications loss during a power recovery; or a spare-part shortage while the same technician is needed elsewhere. The goal is not theatrical catastrophe. It is to reveal hidden couplings before the mission encounters them.

Each scenario should end with a recovery evidence package. The team records what failed, what was known at each decision point, what temporary configuration was used, what reserve was consumed, what debt remains and which gate allows progression. This turns exercises into architecture data rather than anecdotes.

Competency 12 — close the mission with an auditable acceptance matrix

A final acceptance matrix should list the protected functions against the mission phases in which they are required: transit, arrival, early surface, long campaign, contingency and return. For each cell, the evidence identifies nominal capacity, degraded capability, verified recovery path, crew role, critical spares and the decision gate that protects the function. An unresolved cell is not hidden by an impressive overall score; it remains an explicit open risk.

The architecture is ready for approval only when the team can explain why each protected function survives the mission sequence, how uncertainty and common causes are bounded, and what happens when the plan is wrong. This is the essence of the final module: not predicting a perfect Mars campaign, but building a system whose decisions remain intelligible when the campaign stops being perfect.


Integration-board drills — force the dependencies into the open

1. **Dust event plus thermal derating.** Solar output falls while radiator performance also degrades. Show which protected functions lose margin first and how the black-start plan changes.
2. **Water-loop failure with a sick crew specialist.** Rebuild the task and qualification ledger before assuming the nominal repair time is still credible.
3. **Rover battery fault during remote equipment maintenance.** Decide whether to continue the field task, retrieve equipment, protect rescue range or postpone the work.
4. **Communications loss during a power HOLD.** Apply the local authority matrix and record the evidence package that will be transmitted later.
5. **Medical event consumes oxygen and crew time.** Recalculate the protected function horizons without treating medical capability as isolated from ECLSS and workload.
6. **Cannibalisation restores one pump.** Record the new redundancy state and maintenance debt rather than declaring the loop fully recovered.
7. **Critical spare exists in an inaccessible cache.** Treat access time and environment as part of availability, then redesign cache distribution.
8. **Departure vehicle maintenance slips into the return window preparation period.** Protect the return branch by trading local work, crew time and surface objectives explicitly.

Primary bridges for this integration closure: NASA Systems Engineering Handbook, NASA Moon to Mars Architecture, and NASA ECLSS reference.

Mission-command qualification — manage coupled degradation without losing the return

This layer treats the final mission as a living system of protected functions, human qualifications, repair resources, recovery sequences and decision authority. The objective is not to build a larger checklist. It is to make the mission's hidden couplings visible enough that a board can defend what remains possible after several things go wrong at once.



Mission dependency control: find the common causes
Operational context + deterministic dependency map before a contingency is approved

Protected functions and dependencies

```

graph TD
    POWER --- WATER
    POWER --- THERMAL
    THERMAL --- MEDICAL
    MEDICAL --- WATER
    MEDICAL --- RETURN
    THERMAL --- MOBILITY
    MOBILITY --- RETURN
    style MEDICAL fill:#1a3a5a,color:#fff
    style RETURN fill:#8b6914,color:#fff
  
```

Board question
What single loss can simultaneously shorten survival, recovery and return horizons?

Red node = potential common-cause dependency
Ask whether one bus, one operator, one spare class, one software baseline or one location couples the failures.

A mission can fail through a shared dependency even when every subsystem has its own spare. The deterministic network makes common causes visible before the contingency board needs them.

1. Protected functions are the real architecture

Hardware is replaceable in a design; mission functions are what must survive. “Provide breathable atmosphere”, “maintain potable water”, “reject heat”, “preserve medical response”, “move a rescue team” and “keep the departure branch viable” are functional statements. Each can be implemented by several pieces of hardware, procedures and people. Starting from functions prevents the review from confusing a component's presence with the mission's ability to perform the function.

For every protected function, the board should identify the normal chain, degraded chain, detection method, safe state, recovery path and maximum tolerable outage. That last quantity turns a vague priority into a clock. A failure that leaves ten days of margin is managed differently from one that leaves thirty minutes, even if both appear “critical” on a hazard list.

The NASA Systems Engineering Handbook is a useful primary bridge because it treats verification, interfaces and life-cycle thinking as part of engineering rather than after-the-fact documentation.

2. Dependency maps reveal common causes that redundancy diagrams miss

Two pumps are not independent if they share the same power bus, coolant loop, software image, location, operator skill or contamination source. Two communications paths are not independent if one dust event blocks both antennas. Two rescue vehicles are not independent if both require the same charger. Redundancy therefore has to be traced through utilities, controls, environment, maintenance and human operation.

A dependency map should distinguish physical flow, information flow, power, thermal coupling and human qualification. Common-cause candidates are marked explicitly and challenged during review. The practical test is simple: “what one loss can defeat several apparently separate protections at once?” If the answer is a shared bus or one specialist, the architecture has a common-cause vulnerability even if the block diagram shows two boxes.

3. Time-to-loss-of-function orders the control room

During a multi-system event, the loudest alarm should not automatically get the first crew member. The mission needs an estimate of how long each protected function can remain acceptable before a hard limit is crossed. Stored oxygen, thermal inertia, battery energy, medical stabilization time, water reserve, communications autonomy and rover consumables create different horizons.

The board sorts actions by those horizons while considering recovery duration and prerequisite chains. If thermal control will reach a hard limit in two hours but a water issue has two days of protected inventory, the control room may need to stabilize thermal first even if the water alarm looks more dramatic. The horizon must be re-estimated as conditions change; it is a state variable, not a label fixed at the beginning of the emergency.

4. Human workload is a consumable with qualification constraints

Degraded operations often replace automation with manual monitoring, frequent sampling, local valve actions, inspection and logging. That consumes crew time exactly when maintenance and medical workload may also increase. More importantly, not every crew member can perform every task. Workload must therefore be budgeted by qualification, shift, fatigue and backup depth.

A useful ledger lists the degraded task, skill required, minutes per shift, physical location, cognitive burden, backup operator and maximum sustainable duration. This exposes situations where the same electrical specialist, medical officer or ECLSS operator becomes the common cause for several recovery paths. Cross-training adds resilience only when it is current enough to be used under stress.

5. Inventory has states: present, accessible, qualified, available and protected

A spare inside a cargo pallet is not yet a usable spare. It may be physically inaccessible behind other cargo, lack the required tool, be unverified after storage, belong to a different configuration or be reserved for the departure vehicle. Mission inventory therefore needs state labels. “On Mars” is not a sufficient operational status.

During contingency planning, the board should ask where the item is, how long retrieval takes, whether the route is available, whether environmental exposure changed its qualification, and whether consuming it destroys another protected option. Cannibalization must be tracked because removing a part from a dormant system can create hidden debt that matters later.

6. Maintenance debt is a real mission liability

A workaround can restore function while leaving redundancy missing, inspection overdue, software in a temporary configuration or a spare consumed. Calling the mission “recovered” at that moment hides debt. The debt ledger records what is still degraded, how long the temporary state is allowed to continue, which tasks depend on it and what must be done to restore the intended architecture.

This matters over a 500-sol campaign because repeated small deferrals accumulate. A crew can survive several events yet enter the next month with weaker maintenance margin, fewer qualified spares and more manual work. The mission board should therefore review debt as part of the state of the system, not as a list of paperwork to clear when time permits.

7. Medical capability is a chain under a clock

Medical readiness is not the number of medication packages. It depends on detection, diagnosis, trained personnel, sterile or clean procedure capability, monitoring, power, oxygen, communications, evacuation or isolation options and the patient's time window. A medical event can also remove one of the crew's technical specialists, changing the architecture outside medicine.

The control room should therefore treat medical capability as a protected function with dependencies. If the medical officer is also the only person qualified for a critical maintenance task, the staffing model contains a coupling. If the procedure needs power and thermal stability that are already degraded, the medical response cannot be evaluated separately from the engineering event.

8. Communication loss tests delegated authority


Earth cannot be the hidden final step in every decision chain. A Mars surface crew must know which decisions are pre-authorized, which require local technical consensus, which require a commander, and which conditions trigger a conservative safe state while waiting for support. The rule should be written before communication loss, not improvised when latency or outage removes access to specialists.

The mission package should include thresholds for shutting down experiments, isolating equipment, consuming contingency stores, dispatching a rescue vehicle and protecting the departure system. Autonomy here means defined decision authority with evidence requirements, not unlimited discretion.

9. Black-start must restore evidence before discretionary capability

After a major power loss, restart is dangerous because many systems demand current at once and some failures can be re-energized accidentally. A black-start procedure first isolates the failed branch, verifies storage and distribution state, then restores instrumentation and control so the team can observe what happens. Essential thermal and life-support support follows before noncritical loads.

Each step has a measurable acceptance condition. A bus is not “back” because voltage appeared; it must remain stable under the intended load. A water processor is not “recovered” because the pump turns; output quality must be verified. A communications system is not “restored” because one packet passed; the required link performance and routing must be demonstrated.



Black-start recovery: restart without spending the return

Illustrative hardware context + deterministic restart order and return-protection gates

Recovery proof

Restored function ≠ restored redundancy ≠ restored mission margin.

Recovery sequence

- 1 · isolate the failed branch
- 2 · verify safe DC / storage state
- 3 · energize controls + instrumentation
- 4 · restart life-support thermal support
- 5 · add water / comm / medical loads
- 6 · restore discretionary mission loads last

Protected return branch

DO NOT SHED: departure power + avionics + comm + propel

GO restart step verified under measured load

HOLD margin uncertain or common cause unresolved

NO-GO restart would consume protected return capability

Recovery is a sequence, not a switch. The restart order restores evidence and essential functions first while keeping the return branch explicitly protected.

10. The return branch is a protected parallel mission

The departure vehicle, departure power, avionics, communications, propellant conditioning, navigation data and required crew qualifications form a parallel protected branch throughout surface operations. Using those resources for a local contingency can be justified, but the decision must explicitly acknowledge what return capability is being traded.

The practical discipline is to mark return-critical inventory and infrastructure in the same way that emergency reserve is marked in a financial ledger: not invisible, not casually borrowable, and released only by a board that understands the new mission state. A local recovery that quietly consumes the only departure option may turn a solvable surface event into a strategic failure.

11. Failure detection must be independent from the function that failed

If a controller both operates a process and declares itself healthy, some faults can defeat function and detection together. Independent measurements, cross-checks and physical inspection create confidence that recovery is real. The level of independence should match consequence; not every sensor needs a duplicate, but critical functions need a credible way to detect false normal indications.

For example, water production should be separated from water-quality verification. Electrical distribution health can be cross-checked by local voltage/current measurements and downstream behaviour. A software recovery can be checked against external process variables. This prevents the control room from treating a self-reported “nominal” state as proof.

12. Recovery gates should distinguish function, redundancy and mission margin

Three states are often compressed into the word recovered. First, the immediate function works again. Second, the intended redundancy or backup path is restored. Third, the wider mission margin has been rebuilt: reserves replenished, spares replaced, crew workload normalized, temporary configuration removed and return capability restored. These states can occur days apart.

The board should label them separately. A mission may continue limited operations after function restoration while remaining on HOLD for discretionary EVA or science until redundancy and margin are restored. This gives management a vocabulary that does not force every event into “failed” or “nominal”.

13. Logistics must be tested as retrieval, not counting

Mass accounting is necessary but not sufficient. During an emergency the mission needs the correct item, in the correct configuration, at the correct location, with the tool and person needed to use it. Retrieval time can be a mission variable. Cargo stowage, labeling, digital inventory accuracy and physical access routes are therefore part of resilience.

A logistics drill should select a critical spare and trace the full chain: identify it in the database, confirm location, obtain transport, access the container, verify configuration, move it through airlocks if needed, install it, test it and update inventory. If the exercise takes six hours while the protected function has a two-hour horizon, the spare does not close the risk.

14. Crew competence has depth and expiration

Qualification should not be recorded as a permanent yes/no badge. Some skills decay without practice. Others require a second person for verification. Emergency procedures can be familiar in theory but slow under pressure if not rehearsed. The mission should track who is primary, backup and supervised trainee for critical functions, plus when each last practised the task.

Loss of one crew member should be injected into mission exercises specifically to expose single-person dependencies. The objective is not to make everyone an expert in everything; it is to ensure that critical actions have enough depth that one illness, injury or fatigue event does not remove the only executable recovery path.

15. Science must fit inside maintenance reality

Science and exploration are mission objectives, but they consume power, crew time, mobility, samples, communications and maintenance margin. A resilient architecture schedules science against the current health of the system. When maintenance debt grows or protected reserves fall, discretionary activity should be reduced before it consumes recovery options.

This does not mean science is always the first load shed. Some observations are time-critical or can improve safety and resource knowledge. The board should make the trade explicit: expected mission value, resources consumed, reversibility and what protected margin remains afterward.

16. A final acceptance matrix should be adversarial

The final review is stronger when one team is assigned to challenge each assumption, dependency and claimed margin. The mission is not approved because every subsystem lead presents a green slide. It is approved when the integrated architecture survives contradiction: common causes are exposed, degraded states are exercised, critical inventories are physically accessible, return is protected, human workload is credible and unresolved risks have explicit owners.

Acceptance should include evidence dates and expiration where appropriate. A test performed months before launch may no longer represent a software or hardware configuration that changed. A spare count can change after maintenance. Crew qualification can lapse. Mission closure is therefore a controlled snapshot with change rules, not a one-time ceremony.

17. Coupled-failure campaign — twelve integration-board situations

1. **Power loss plus water contamination alarm.** Prioritize by time-to-loss-of-function, preserve diagnostic power and verify water quality independently before assuming the storage is unusable.
2. **One ECLSS specialist becomes medically unavailable.** Recalculate staffing depth and manual degraded tasks; do not assume the remaining crew can absorb the workload indefinitely.
3. **A rover failure removes the only fast rescue route.** Update EVA radius and refuge strategy immediately rather than waiting for the next EVA plan.
4. **Black-start succeeds electrically but water output remains unqualified.** Mark power function restored while water function remains on HOLD.
5. **A cannibalized valve restores habitat cooling.** Record the donor system's new state and the maintenance debt created by the recovery.
6. **Communication outage coincides with a cargo-fire alarm.** Apply pre-delegated local authority and safe-state thresholds; do not wait for Earth if the hazard clock is shorter than the communications horizon.
7. **Medical procedure needs the same power branch reserved for black-start.** Make the coupling explicit and compare the medical time window with the system-recovery clock.
8. **Inventory database says a spare exists but the container is blocked by cargo.** Treat retrieval time as part of availability and redesign stowage for the next cycle.
9. **Surface survival can be extended by using departure batteries.** Require an explicit board decision because return readiness is being traded for local horizon.
10. **Two redundant controllers share one software defect.** Redundancy is not independence; use diverse rollback or external safe control if available.
11. **Maintenance debt accumulates after three “successful” workarounds.** Reduce discretionary operations until redundancy, spares and crew workload recover.
12. **A new crew inherits a system with undocumented temporary configuration.** Hold high-risk operations until configuration is reconciled and the evidence pack is rebuilt.

18. Final handover package — what the next crew must be able to reproduce

The handover should contain the protected-function map, dependency network, current inventory states, maintenance debt, open hazards, configuration baselines, crew qualification depth, emergency authority rules, black-start sequence, rescue geometry, medical contingencies, return-readiness state and evidence behind every current GO/HOLD/NO-GO disposition. The incoming crew should be able to challenge the architecture without needing the outgoing crew's memory.

The goal is institutional continuity. A Mars settlement or long expedition cannot depend on one heroic team understanding a pile of undocumented exceptions. The architecture becomes durable when the reasoning, evidence and recovery options survive personnel change.

Primary-source bridges: NASA Systems Engineering Handbook, NASA Moon to Mars Architecture components, NASA ECLSS reference, and NASA human spaceflight resources.

Integrated mission-command extension — authority, configuration and resilience across a long campaign

The architecture is not finished when every subsystem has a diagram and every nominal budget closes. Long-duration mission command must remain coherent when people are tired, parts are unavailable, software changes, cargo is late, communications disappear and several apparently independent systems fail through one shared dependency. This extension turns those interactions into reviewable command problems.

19. Incident command must follow the hazard clock

Authority should be designed before the emergency. A cabin fire, toxic release, pressure loss, medical emergency or power collapse may evolve faster than Earth can participate. The crew therefore needs a predeclared incident-command model: who takes command, who protects life-support continuity, who owns medical decisions, who controls electrical isolation, who communicates, who records configuration changes and who can stop an unsafe action.

The command structure should change with the phase. Immediate response favors speed and clear authority. Stabilization favors independent verification and disciplined configuration control. Recovery favors system owners, maintenance expertise and evidence. Return-to-normal requires explicit acceptance criteria so that the temporary emergency organization does not silently become the permanent operating mode.

20. Change control prevents yesterday's fix from becoming tomorrow's common cause

A software patch, rewired power feed, bypassed valve or temporary procedure can restore a function while creating a latent risk elsewhere. The mission therefore needs configuration control proportionate to consequence. Each temporary change should have an owner, purpose, affected interfaces, verification evidence, expiration or review condition, rollback path and documentation visible to the next shift.

Emergency change control is not bureaucracy for its own sake. It is memory for a system whose operators rotate and whose failures can interact. The rule is simple: act quickly when the hazard clock requires it, but record the configuration immediately enough that the next decision is based on the spacecraft or habitat that actually exists, not the one shown in an old diagram.

21. Software redundancy is not independence when the code is shared

Two computers running the same flawed build can fail together. Two controllers using the same corrupted input can make the same wrong decision. A mission board should therefore distinguish hardware redundancy from functional independence. Diversity may come from separate sensing, simpler backup logic, manual control, a validated previous software version or an external monitoring path.

The important question is not “how many computers do we have?” but “what failures can defeat all of them at once?”. This common-cause review belongs beside redundancy claims in power, ECLSS, navigation, medical support and mission software.

22. Local manufacturing must include qualification, not merely fabrication

A settlement that can print or machine a replacement part is more resilient than one that depends entirely on Earth, but fabrication capability alone does not prove the part is safe. The team must control feedstock identity, dimensional inspection, material properties, process history, cleaning, compatibility and functional test appropriate to the consequence of failure.

For low-consequence items, visual inspection and fit checks may be enough. For pressure boundaries, breathing-gas components, structural links or critical electrical hardware, the qualification burden is much higher. A local workshop therefore needs test capability, reference artifacts, calibration discipline and a clear rule for when a locally produced item is temporary, restricted or fully accepted.

23. Spare strategy should be organized by protected function

Counting part numbers can hide fragility. One spare pump may serve several loops but only if interfaces and operating conditions are compatible. One seal kit may protect multiple functions, while a unique electronic module may be a single point of long-lead failure. The spare strategy should therefore map parts to protected functions, failure frequency, repair time, diagnostic confidence and alternatives such as cannibalization or local manufacture.

The best inventory is not necessarily the largest. It is the inventory that preserves mission functions across plausible failure combinations while remaining accessible and qualified. This is why stowage location, retrieval equipment and documentation belong in the same review as quantity.

24. Risk registers must not multiply probabilities that are not independent

Campaign risk is often summarized with numbers, but common causes can make apparently separate events correlated. Dust can degrade power, thermal rejection, visibility and mechanisms at once. A crew illness can remove both maintenance expertise and medical capacity. One software defect can affect redundant controllers. Treating such events as independent can produce falsely reassuring combined probabilities.

The review should identify shared causes and use scenarios or conditional reasoning when independence is not defensible. Quantitative risk models are useful, but only when their assumptions remain visible. The board should be able to point to the common-cause groups and explain how the architecture breaks them.

25. Rare failure preparation must fit inside routine workload

A procedure that requires twelve hours of uninterrupted expert work may be unrealistic during the failure it is supposed to solve. Recovery plans must include staffing, protective equipment, access time, tool preparation, communication, rest, monitoring and concurrent essential operations. A technically correct repair can still be operationally impossible if it consumes the only qualified operator for another critical function.

Exercise design should therefore measure more than task completion. Record operator-minutes, specialist bottlenecks, handovers, error opportunities, required permits, consumables used and functions temporarily left unattended. Those measurements turn “we have a procedure” into evidence that the procedure can be executed under campaign conditions.

26. Handover is a safety-critical interface between people

Long missions run across shifts, crew rotations and eventually generations of operators. Handover quality therefore deserves the same systems-engineering discipline as a hardware interface. The outgoing team should transmit configuration state, active hazards, degraded functions, open maintenance, consumable margins, temporary procedures, upcoming decision gates and the evidence behind unresolved anomalies.

A strong handover is testable: the incoming operator can explain the current risk picture, identify what would cause escalation and reproduce the key evidence without relying on undocumented conversation. If that test fails, the mission has an information single point of failure.

27. Return capability must survive local success

Surface operations constantly tempt the crew to borrow from return assets: batteries, propellant reserves, spare hardware, medical stock, communications equipment or trained personnel. Sometimes that trade is rational, but it must be explicit. Every proposed use should state what return function is degraded, how it can be restored, how long restoration takes and what event would make restoration impossible.

The command team should maintain a protected-return ledger separate from general inventory. A settlement can be locally stable yet strategically trapped if departure hardware, ascent access or rendezvous support has been quietly consumed.

28. The seventy-two-hour integrated incident

Use this capstone as a command exercise, not a reading quiz. At the start, one power string trips during a dust event. Water recovery output is uncertain, not confirmed failed. One rover is outside its preferred maintenance interval. A crew member develops symptoms that may reduce specialist availability. Communications with Earth are intermittent. A software rollback is available but would remove a recent fault-detection feature. The return vehicle is healthy, although one of its protected energy reserves could temporarily stabilize the habitat.

During the first hour, identify protected functions, time-to-loss-of-function and the local authority needed before Earth can answer. During the first shift, establish independent evidence, isolate common causes and create a configuration log. During the next twenty-four hours, restore sufficient redundancy without exhausting the crew. By seventy-two hours, decide which temporary configurations can remain, which maintenance debt must be paid immediately, whether science stays reduced and whether any return resource borrowed during stabilization has been fully restored.

The exercise is successful only if the team can defend the sequence of decisions. “Everything is green at the end” is not enough. The evidence should show that the crew never lost sight of survival, recovery, workload, common cause, configuration truth and the protected return branch.

29. Final mission-command board — questions that must have explicit answers

1. Which functions are protected, and what are their current horizons?
2. Which apparent redundancies still share power, software, crew, cooling, location or procedures?
3. Which temporary changes are active, who owns them and when must they be reviewed?
4. Which spares are physically accessible and qualified rather than merely listed?
5. Which local-manufacturing capabilities can produce a part, and which can actually qualify it?
6. Which recovery procedure creates the largest specialist bottleneck?
7. What authority is delegated locally during communication loss?
8. What evidence independently confirms that ECLSS and power have recovered?
9. Which maintenance debt is acceptable for another sol, and which threatens a protected function?
10. Which return resources have been borrowed, and what is the verified restoration path?
11. What is the current medical capability if the most qualified crew member becomes the patient?
12. What must the next shift know to reproduce the current GO/HOLD/NO-GO state?

Primary-source bridge: NASA Systems Engineering Handbook, NASA Moon to Mars Architecture components, and NASA ECLSS reference.

Mission-closure calculation extension — six protected operational horizons

These six mini-lessons strengthen the final architecture course with explicit power, oxygen, food, maintenance, spares and communications-autonomy calculations tied to GO/HOLD/NO-GO decisions.

Critical-load energy requirement

$$E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$$

1 — Concrete question

For **Critical-load energy requirement**, how does $E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$ inform closing the energy needed to hold a protected electrical function for a defined duration and the operational choice “Keep critical-load energy closure separate from total settlement energy so discretionary loads cannot consume survival reserve.”?

2 — Intuition without symbols

Intuition. A critical device consumes more energy when it draws more power or must remain active longer. Mission planning therefore turns an emergency duration into an energy reserve before deciding whether a battery or generator can support it.

3 — Quantities first

P_{crit} is protected load power; t_{support} required support duration; E_{req} energy delivered to the load.

4 — Formula

$$E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$$

5 — Read aloud

“E required equals P critical times t support.”

6 — Symbols

Symbol map for Critical-load energy requirement. P_{crit} is protected load power; t_{support} required support duration; E_{req} energy delivered to the load.

7 — Pronunciation

Pronunciation. Say $E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$. For *Critical-load energy requirement*, use the step-three names tied to closing the energy needed to hold a protected electrical function for a defined duration. Speak each *Critical-load energy requirement* unit with the quantity it measures.

8 — Units

$\text{kW} \times \text{h} = \text{kWh}$

9 — Convention

Convention. For *Critical-load energy requirement*, keep closing the energy needed to hold a protected electrical function for a defined duration on one declared boundary. Apply $E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$ under that convention. Storage capacity must also include conversion efficiency, degradation, temperature and protected reserve.

10 — Why this operation

Why this operation. $E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$ answers the *Critical-load energy requirement* question because it represents closing the energy needed to hold a protected electrical function for a defined duration. In this case it yields: The critical load requires 72 kWh delivered over 18 hours.

11 — Assumptions

Assumptions. Treat the *Critical-load energy requirement* values as one teaching case. For closing the energy needed to hold a protected electrical function for a defined duration, keep a single physical or operational boundary. Storage capacity must also include conversion efficiency, degradation, temperature and protected reserve.

12 — Unit check

Unit check. Reduce $E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$ for *Critical-load energy requirement*. The required dimension is $\text{kW} \times \text{h} = \text{kWh}$. A different dimension invalidates “The critical load requires 72 kWh delivered over 18 hours.”.

13 — Numerical case

$P_{\text{crit}} = 4.0 \text{ kW}$

$t_{\text{support}} = 18 \text{ h}$

$E_{\text{req}} = 4.0 \times 18 = 72 \text{ kWh}$

14 — Why each operation

Why each operation. For *Critical-load energy requirement*, substitute $P_{crit} = 4.0 \text{ kW}$; $t_{support} = 18 \text{ h}$; $E_{req} = 4.0 \times 18 = 72 \text{ kWh}$ into $E_{req} = P_{crit} \times t_{support}$. Then verify the independent statement " $72/18=4.0 \text{ kW}$ ".

15 — Algebra check

Algebra check. Reverse $E_{req} = P_{crit} \times t_{support}$ for *Critical-load energy requirement* using " $72/18=4.0 \text{ kW}$ ". The recovered input should follow "A 25% longer support horizon raises required delivered energy by 25%.". If not, recheck units and boundaries.

16 — Mental estimate

Mental estimate. Round the dominant inputs for *Critical-load energy requirement*. Compare that rough scale with "The critical load requires 72 kWh delivered over 18 hours.". If they diverge sharply, inspect $E_{req} = P_{crit} \times t_{support}$ for units, signs or boundaries.

17 — Interpretation

Interpretation. For *Critical-load energy requirement*, The critical load requires 72 kWh delivered over 18 hours. Operationally: Keep critical-load energy closure separate from total settlement energy so discretionary loads cannot consume survival reserve. The interpretation remains limited by "Storage capacity must also include conversion efficiency, degradation, temperature and protected reserve.".

18 — What it does not prove

What it does not prove. *Critical-load energy requirement* cannot support claims outside closing the energy needed to hold a protected electrical function for a defined duration. Storage capacity must also include conversion efficiency, degradation, temperature and protected reserve. Use the result only to justify: Keep critical-load energy closure separate from total settlement energy so discretionary loads cannot consume survival reserve.

19 — Sensitivity or limit case

A 25% longer support horizon raises required delivered energy by 25%.

20 — Practice

Guided exercise — Critical-load energy requirement. A 3.5 kW load must survive 24 h. Find delivered energy.

Guided correction — Critical-load energy requirement

1. $E = 3.5 \times 24 = 84$ kWh.
2. Then size storage above 84 kWh for losses and reserve.

Autonomous exercise — Critical-load energy requirement. Build a second case from “A 25% longer support horizon raises required delivered energy by 25%.”. Re-evaluate $E_{\text{req}} = P_{\text{crit}} \times t_{\text{support}}$. Name the changed input. Decide whether “Keep critical-load energy closure separate from total settlement energy so discretionary loads cannot consume survival reserve.” still follows.

Autonomous correction — Critical-load energy requirement

For *Critical-load energy requirement*, state the altered case. Preserve $\text{kW} \times \text{h} = \text{kWh}$. Match the direction in “A 25% longer support horizon raises required delivered energy by 25%.”. Respect “Storage capacity must also include conversion efficiency, degradation, temperature and protected reserve.”. Finish by retaining or revising: Keep critical-load energy closure separate from total settlement energy so discretionary loads cannot consume survival reserve.

21 — Mission decision

Keep critical-load energy closure separate from total settlement energy so discretionary loads cannot consume survival reserve.

Crew oxygen reserve horizon

$$t_{\text{O}_2} = M_{\text{O}_2\text{usable}} / (N_{\text{crew}} \times q_{\text{O}_2})$$

1 — Concrete question

For **Crew oxygen reserve horizon**, how does $t_{\text{O}_2} = M_{\text{O}_2\text{usable}} / (N_{\text{crew}} \times q_{\text{O}_2})$ inform estimating how long an accessible oxygen reserve can support the crew at a stated planning consumption and the operational choice “Use oxygen horizon in the same time-to-loss board as CO₂ removal and pressure integrity so the earliest clock drives action.”?

2 — Intuition without symbols

Intuition. An oxygen reserve lasts only as long as the usable stock can cover the combined breathing demand of the crew. More people or greater consumption shortens the horizon; more protected oxygen extends it.

3 — Quantities first

M_O2_usable is qualified accessible oxygen mass; N_crew crew count; q_O2 planning consumption per person per day; t_O2 is horizon.

4 — Formula

$$t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$$

5 — Read aloud

"t O two equals usable O two mass divided by crew count times per-person O two rate."

6 — Symbols

Symbol map for Crew oxygen reserve horizon. M_O2_usable is qualified accessible oxygen mass; N_crew crew count; q_O2 planning consumption per person per day; t_O2 is horizon.

7 — Pronunciation

Pronunciation. Say $t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$. For *Crew oxygen reserve horizon*, use the step-three names tied to estimating how long an accessible oxygen reserve can support the crew at a stated planning consumption. Speak each *Crew oxygen reserve horizon* unit with the quantity it measures.

8 — Units

$$\text{kg} / (\text{persons} \times \text{kg/person/day}) = \text{day}$$

9 — Convention

Convention. For *Crew oxygen reserve horizon*, keep estimating how long an accessible oxygen reserve can support the crew at a stated planning consumption on one declared boundary. Apply $t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$ under that convention. Metabolic demand varies and oxygen can be constrained by delivery hardware, pressure and fire strategy; this is an inventory horizon.

10 — Why this operation

Why this operation. $t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$ answers the *Crew oxygen reserve horizon* question because it represents estimating how long an accessible oxygen reserve can support the crew at a stated planning consumption. In this case it yields: The reserve supports about 7.1 days at the stated planning rate.

11 — Assumptions

Assumptions. Treat the *Crew oxygen reserve horizon* values as one teaching case. For estimating how long an accessible oxygen reserve can support the crew at a stated planning consumption, keep a single physical or operational boundary. Metabolic demand varies and oxygen can be constrained by delivery hardware, pressure and fire strategy; this is an inventory horizon.

12 — Unit check

Unit check. Reduce $t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$ for *Crew oxygen reserve horizon*. The required dimension is $kg / (persons \times kg/person/day) = day$. A different dimension invalidates “The reserve supports about 7.1 days at the stated planning rate.”.

13 — Numerical case

$M_{O2_usable} = 36 \text{ kg}$

$N_{crew} = 6$

$q_{O2} = 0.84 \text{ kg/person/day}$

$t_{O2} = 36 / (6 \times 0.84) \approx 7.14 \text{ days}$

14 — Why each operation

Why each operation. For *Crew oxygen reserve horizon*, substitute $M_{O2_usable} = 36 \text{ kg}$; $N_{crew} = 6$; $q_{O2} = 0.84 \text{ kg/person/day}$; $t_{O2} = 36 / (6 \times 0.84) \approx 7.14 \text{ days}$ into $t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$. Then verify the independent statement “ $6 \times 0.84 \times 7.14 \approx 36 \text{ kg}$ ”.

15 — Algebra check

Algebra check. Reverse $t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$ for *Crew oxygen reserve horizon* using “ $6 \times 0.84 \times 7.14 \approx 36 \text{ kg}$ ”. The recovered input should follow “Adding two crew with the same reserve reduces horizon by 25%.”. If not, recheck units and boundaries.

16 — Mental estimate

Mental estimate. Round the dominant inputs for *Crew oxygen reserve horizon*. Compare that rough scale with “The reserve supports about 7.1 days at the stated planning rate.”. If they diverge sharply, inspect $t_{O2} = M_{O2_usable} / (N_{crew} \times q_{O2})$ for units, signs or boundaries.

17 — Interpretation

Interpretation. For *Crew oxygen reserve horizon*, The reserve supports about 7.1 days at the stated planning rate. Operationally: Use oxygen horizon in the same time-to-loss board as CO₂ removal and pressure integrity so the earliest clock drives action. The interpretation remains limited by “Metabolic demand varies and oxygen can be constrained by delivery hardware, pressure and fire strategy; this is an inventory horizon.”.

18 — What it does not prove

What it does not prove. *Crew oxygen reserve horizon* cannot support claims outside estimating how long an accessible oxygen reserve can support the crew at a stated planning consumption. Metabolic demand varies and oxygen can be constrained by delivery hardware, pressure and fire strategy; this is an inventory horizon. Use the result only to justify: Use oxygen horizon in the same time-to-loss board as CO₂ removal and pressure integrity so the earliest clock drives action.

19 — Sensitivity or limit case

Adding two crew with the same reserve reduces horizon by 25%.

20 — Practice

Guided exercise — Crew oxygen reserve horizon. $M=50$ kg, $N=5$, $q=0.84$ kg/person/day. Find horizon.

Guided correction — Crew oxygen reserve horizon

1. $t=50/(4.2)\approx 11.9$ days.
2. Keep inaccessible and emergency-only inventory out of routine usable mass.

Autonomous exercise — Crew oxygen reserve horizon. Build a second case from “Adding two crew with the same reserve reduces horizon by 25%.”. Re-evaluate $t_{O_2} = M_{O_2_usable} / (N_{crew} \times q_{O_2})$. Name the changed input. Decide whether “Use oxygen horizon in the same time-to-loss board as CO₂ removal and pressure integrity so the earliest clock drives action.” still follows.

Autonomous correction — Crew oxygen reserve horizon

For *Crew oxygen reserve horizon*, state the altered case. Preserve $kg / (persons \times kg/person/day) = day$. Match the direction in “Adding two crew with the same reserve reduces horizon by 25%.”. Respect “Metabolic demand varies and oxygen can be constrained by delivery hardware, pressure and fire strategy; this is an inventory horizon.”. Finish by retaining or revising: Use oxygen horizon in the same time-to-loss board as CO₂ removal and pressure integrity so the earliest clock drives action.

21 — Mission decision

Use oxygen horizon in the same time-to-loss board as CO₂ removal and pressure integrity so the earliest clock drives action.

Food stock horizon with protected reserve

$$t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$$

1 — Concrete question

For **Food stock horizon with protected reserve**, how does $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$ inform separating routine food stock from an emergency reserve that cannot be casually consumed and the operational choice “Treat reserve release as a command decision with a replenishment/recovery plan, not as hidden schedule margin.”?

2 — Intuition without symbols

Intuition. Food autonomy is a stock-and-demand problem. The useful food mass is what remains after protecting the reserve, and that usable amount must cover every crewmember for the planned duration.

3 — Quantities first

M_{food} is total edible stock; M_{reserve} protected mass; N_{crew} crew size; q_{food} planning consumption mass per person per day.

4 — Formula

$$t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$$

5 — Read aloud

“ t_{food} equals total food minus reserve divided by crew count times food rate.”

6 — Symbols

Symbol map for Food stock horizon with protected reserve. M_{food} is total edible stock; M_{reserve} protected mass; N_{crew} crew size; q_{food} planning consumption mass per person per day.

7 — Pronunciation

Pronunciation. Say $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$. For *Food stock horizon with protected reserve*, use the step-three names tied to separating routine food stock from an emergency reserve that cannot be casually consumed. Speak each *Food stock horizon with protected reserve* unit with the quantity it measures.

8 — Units

$$\text{kg} / (\text{persons} \times \text{kg/person/day}) = \text{day}$$

9 — Convention

Convention. For *Food stock horizon with protected reserve*, keep separating routine food stock from an emergency reserve that cannot be casually consumed on one declared boundary. Apply $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$ under that convention. Food mass does not capture nutrition, menu compatibility, spoilage or packaging losses; reserve policy must state what triggers release.

10 — Why this operation

Why this operation. $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$ answers the *Food stock horizon with protected reserve* question because it represents separating routine food stock from an emergency reserve that cannot be casually consumed. In this case it yields: Routine usable stock covers 75 days while 60 kg remains protected.

11 — Assumptions

Assumptions. Treat the *Food stock horizon with protected reserve* values as one teaching case. For separating routine food stock from an emergency reserve that cannot be casually consumed, keep a single physical or operational boundary. Food mass does not capture nutrition, menu compatibility, spoilage or packaging losses; reserve policy must state what triggers release.

12 — Unit check

Unit check. Reduce $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$ for *Food stock horizon with protected reserve*. The required dimension is $\text{kg} / (\text{persons} \times \text{kg/person/day}) = \text{day}$. A different dimension invalidates “Routine usable stock covers 75 days while 60 kg remains protected.”.

13 — Numerical case

$$M_{\text{food}} = 420 \text{ kg}$$

$$M_{\text{reserve}} = 60 \text{ kg}$$

$$N_{\text{crew}} = 6$$

$$q_{\text{food}} = 0.80 \text{ kg/person/day}$$

$$t_{\text{food}} = (420 - 60) / (6 \times 0.80) = 360 / 4.8 = 75 \text{ days}$$

14 — Why each operation

Why each operation. For *Food stock horizon with protected reserve*, substitute $M_{\text{food}} = 420 \text{ kg}$; $M_{\text{reserve}} = 60 \text{ kg}$; $N_{\text{crew}} = 6$; $q_{\text{food}} = 0.80 \text{ kg/person/day}$; $t_{\text{food}} = (420 - 60) / (6 \times 0.80) = 360 / 4.8 = 75 \text{ days}$ into $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$. Then verify the independent statement “ $75 \times 4.8 = 360 \text{ kg}$ routine mass”.

15 — Algebra check

Algebra check. Reverse $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$ for *Food stock horizon with protected reserve* using “75×4.8=360 kg routine mass”. The recovered input should follow “If crew consumption rises 10%, horizon falls from 75 to about 68.2 days.”. If not, recheck units and boundaries.

16 — Mental estimate

Mental estimate. Round the dominant inputs for *Food stock horizon with protected reserve*. Compare that rough scale with “Routine usable stock covers 75 days while 60 kg remains protected.”. If they diverge sharply, inspect $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$ for units, signs or boundaries.

17 — Interpretation

Interpretation. For *Food stock horizon with protected reserve*, Routine usable stock covers 75 days while 60 kg remains protected. Operationally: Treat reserve release as a command decision with a replenishment/recovery plan, not as hidden schedule margin. The interpretation remains limited by “Food mass does not capture nutrition, menu compatibility, spoilage or packaging losses; reserve policy must state what triggers release.”.

18 — What it does not prove

What it does not prove. *Food stock horizon with protected reserve* cannot support claims outside separating routine food stock from an emergency reserve that cannot be casually consumed. Food mass does not capture nutrition, menu compatibility, spoilage or packaging losses; reserve policy must state what triggers release. Use the result only to justify: Treat reserve release as a command decision with a replenishment/recovery plan, not as hidden schedule margin.

19 — Sensitivity or limit case

If crew consumption rises 10%, horizon falls from 75 to about 68.2 days.

20 — Practice

Guided exercise — Food stock horizon with protected reserve. $M=300$ kg, reserve=40 kg, $N=5$, $q=0.75$ kg/person/day. Find routine horizon.

Guided correction — Food stock horizon with protected reserve

1. usable=260 kg; rate=3.75 kg/day; $t \approx 69.3$ days.
2. Keep the 40 kg reserve outside the routine schedule.

Autonomous exercise — Food stock horizon with protected reserve. Build a second case from “If crew consumption rises 10%, horizon falls from 75 to about 68.2 days.”. Re-evaluate $t_{\text{food}} = (M_{\text{food}} - M_{\text{reserve}}) / (N_{\text{crew}} \times q_{\text{food}})$. Name the changed input. Decide whether “Treat reserve release as a command decision with a replenishment/recovery plan, not as hidden schedule margin.” still follows.

Autonomous correction — Food stock horizon with protected reserve

For *Food stock horizon with protected reserve*, state the altered case. Preserve $\text{kg} / (\text{persons} \times \text{kg/person/day}) = \text{day}$. Match the direction in “If crew consumption rises 10%, horizon falls from 75 to about 68.2 days.”. Respect “Food mass does not capture nutrition, menu compatibility, spoilage or packaging losses; reserve policy must state what triggers release.”. Finish by retaining or revising: Treat reserve release as a command decision with a replenishment/recovery plan, not as hidden schedule margin.

21 — Mission decision

Treat reserve release as a command decision with a replenishment/recovery plan, not as hidden schedule margin.

Maintenance workload utilization

$$U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$$

1 — Concrete question

For **Maintenance workload utilization**, how does $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$ inform checking whether qualified labor hours can absorb expected maintenance demand and the operational choice “Protect maintenance headroom so a single failure does not force deferral of another safety-critical task.”?

2 — Intuition without symbols

Intuition. Maintenance demand competes for a finite pool of qualified human time. When required work approaches or exceeds the skilled hours actually available, deferred maintenance and fatigue become mission risks rather than scheduling inconveniences.

3 — Quantities first

H_required is maintenance labor hours needed in a period; H_qualified is qualified labor hours actually available after duty/rest and other protected work; U_maint is utilization.

4 — Formula

$$U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$$

5 — Read aloud

“U maintenance equals required hours divided by qualified hours.”

6 — Symbols

Symbol map for Maintenance workload utilization. H_required is maintenance labor hours needed in a period; H_qualified is qualified labor hours actually available after duty/rest and other protected work; U_maint is utilization.

7 — Pronunciation

Pronunciation. Say $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$. For *Maintenance workload utilization*, use the step-three names tied to checking whether qualified labor hours can absorb expected maintenance demand. Speak each *Maintenance workload utilization* unit with the quantity it measures.

8 — Units

h/h = dimensionless

9 — Convention

Convention. For *Maintenance workload utilization*, keep checking whether qualified labor hours can absorb expected maintenance demand on one declared boundary. Apply $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$ under that convention. Hours are not interchangeable if only some crew are qualified for specific tasks; fatigue and task concurrency also matter.

10 — Why this operation

Why this operation. $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$ answers the *Maintenance workload utilization* question because it represents checking whether qualified labor hours can absorb expected maintenance demand. In this case it yields: Maintenance uses 80% of available qualified capacity, leaving 24 h/week before further contingency work.

11 — Assumptions

Assumptions. Treat the *Maintenance workload utilization* values as one teaching case. For checking whether qualified labor hours can absorb expected maintenance demand, keep a single physical or operational boundary. Hours are not interchangeable if only some crew are qualified for specific tasks; fatigue and task concurrency also matter.

12 — Unit check

Unit check. Reduce $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$ for *Maintenance workload utilization*. The required dimension is h/h = dimensionless. A different dimension invalidates “Maintenance uses 80% of available qualified capacity, leaving 24 h/week before further contingency work.”.

13 — Numerical case

$H_{\text{required}} = 96 \text{ h/week}$

$H_{\text{qualified}} = 120 \text{ h/week}$

$U_{\text{maint}} = 96/120 = 0.80 = 80\%$

14 — Why each operation

Why each operation. For *Maintenance workload utilization*, substitute $H_{\text{required}} = 96 \text{ h/week}$; $H_{\text{qualified}} = 120 \text{ h/week}$; $U_{\text{maint}} = 96/120 = 0.80 = 80\%$ into $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$. Then verify the independent statement “ $0.80 \times 120 = 96 \text{ h}$ ”.

15 — Algebra check

Algebra check. Reverse $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$ for *Maintenance workload utilization* using “ $0.80 \times 120 = 96 \text{ h}$ ”. The recovered input should follow “An extra 18 h failure pushes utilization to $114/120 = 95\%$.”. If not, recheck units and boundaries.

16 — Mental estimate

Mental estimate. Round the dominant inputs for *Maintenance workload utilization*. Compare that rough scale with “Maintenance uses 80% of available qualified capacity, leaving 24 h/week before further contingency work.”. If they diverge sharply, inspect $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$ for units, signs or boundaries.

17 — Interpretation

Interpretation. For *Maintenance workload utilization*, Maintenance uses 80% of available qualified capacity, leaving 24 h/week before further contingency work. Operationally: Protect maintenance headroom so a single failure does not force deferral of another safety-critical task. The interpretation remains limited by “Hours are not interchangeable if only some crew are qualified for specific tasks; fatigue and task concurrency also matter.”.

18 — What it does not prove

What it does not prove. *Maintenance workload utilization* cannot support claims outside checking whether qualified labor hours can absorb expected maintenance demand. Hours are not interchangeable if only some crew are qualified for specific tasks; fatigue and task concurrency also matter. Use the result only to justify: Protect maintenance headroom so a single failure does not force deferral of another safety-critical task.

19 — Sensitivity or limit case

An extra 18 h failure pushes utilization to $114/120=95\%$.

20 — Practice

Guided exercise — Maintenance workload utilization. Required=88 h, qualified capacity=100 h. Find utilization.

Guided correction — Maintenance workload utilization

1. $U=88\%$.
2. Identify skill bottlenecks before treating the remaining 12 h as fungible.

Autonomous exercise — Maintenance workload utilization. Build a second case from “An extra 18 h failure pushes utilization to $114/120=95\%$.”. Re-evaluate $U_{\text{maint}} = H_{\text{required}} / H_{\text{qualified}}$. Name the changed input. Decide whether “Protect maintenance headroom so a single failure does not force deferral of another safety-critical task.” still follows.

Autonomous correction — Maintenance workload utilization

For *Maintenance workload utilization*, state the altered case. Preserve $h/h =$ dimensionless. Match the direction in “An extra 18 h failure pushes utilization to $114/120=95\%$.”. Respect “Hours are not interchangeable if only some crew are qualified for specific tasks; fatigue and task concurrency also matter.”. Finish by retaining or revising: Protect maintenance headroom so a single failure does not force deferral of another safety-critical task.

21 — Mission decision

Protect maintenance headroom so a single failure does not force deferral of another safety-critical task.

Critical spare coverage ratio

$$C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$$

1 — Concrete question

For **Critical spare coverage ratio**, how does $C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$ inform testing whether physically accessible qualified spares cover a defined contingency case and the operational choice “Count only spares that are serviceable, reachable and configuration-compatible in the acceptance matrix.”?

2 — Intuition without symbols

Intuition. A spare inventory is useful only when it can cover the failures the scenario actually asks the crew to survive. Counting serviceable replacement units against the protected failure case reveals whether the inventory is adequate or merely comforting.

3 — Quantities first

$N_{\text{serviceable}}$ is accessible, compatible, qualified spare count; $N_{\text{required_case}}$ is units needed for the contingency; C_{spare} is coverage ratio.

4 — Formula

$$C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$$

5 — Read aloud

“C spare equals serviceable spare count divided by required case count.”

6 — Symbols

Symbol map for Critical spare coverage ratio. $N_{\text{serviceable}}$ is accessible, compatible, qualified spare count; $N_{\text{required_case}}$ is units needed for the contingency; C_{spare} is coverage ratio.

7 — Pronunciation

Pronunciation. Say $C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$. For *Critical spare coverage ratio*, use the step-three names tied to testing whether physically accessible qualified spares cover a defined contingency case. Speak each *Critical spare coverage ratio* unit with the quantity it measures.

8 — Units

units/units = dimensionless

9 — Convention

Convention. For *Critical spare coverage ratio*, keep testing whether physically accessible qualified spares cover a defined contingency case on one declared boundary. Apply $C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$ under that convention. A ratio above one does not protect against common-cause defects, wrong configuration or inaccessible storage.

10 — Why this operation

Why this operation. $C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$ answers the *Critical spare coverage ratio* question because it represents testing whether physically accessible qualified spares cover a defined contingency case. In this case it yields: Coverage is 1.5 times the defined contingency requirement.

11 — Assumptions

Assumptions. Treat the *Critical spare coverage ratio* values as one teaching case. For testing whether physically accessible qualified spares cover a defined contingency case, keep a single physical or operational boundary. A ratio above one does not protect against common-cause defects, wrong configuration or inaccessible storage.

12 — Unit check

Unit check. Reduce $C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$ for *Critical spare coverage ratio*. The required dimension is $\text{units}/\text{units} = \text{dimensionless}$. A different dimension invalidates "Coverage is 1.5 times the defined contingency requirement."

13 — Numerical case

$$N_{\text{serviceable}} = 3$$

$$N_{\text{required_case}} = 2$$

$$C_{\text{spare}} = 3/2 = 1.5$$

14 — Why each operation

Why each operation. For *Critical spare coverage ratio*, substitute $N_{\text{serviceable}} = 3$; $N_{\text{required_case}} = 2$; $C_{\text{spare}} = 3/2 = 1.5$ into $C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$. Then verify the independent statement "2 required $\times 1.5 = 3$ serviceable spares".

15 — Algebra check

Algebra check. Reverse $C_{\text{spare}} = N_{\text{serviceable}} / N_{\text{required_case}}$ for *Critical spare coverage ratio* using "2 required $\times 1.5 = 3$ serviceable spares". The recovered input should follow "One quarantined spare reduces coverage from 1.5 to 1.0 in this case.". If not, recheck units and boundaries.

16 — Mental estimate

Mental estimate. Round the dominant inputs for *Critical spare coverage ratio*. Compare that rough scale with “Coverage is 1.5 times the defined contingency requirement.”. If they diverge sharply, inspect $C_spare = N_serviceable / N_required_case$ for units, signs or boundaries.

17 — Interpretation

Interpretation. For *Critical spare coverage ratio*, Coverage is 1.5 times the defined contingency requirement. Operationally: Count only spares that are serviceable, reachable and configuration-compatible in the acceptance matrix. The interpretation remains limited by “A ratio above one does not protect against common-cause defects, wrong configuration or inaccessible storage.”.

18 — What it does not prove

What it does not prove. *Critical spare coverage ratio* cannot support claims outside testing whether physically accessible qualified spares cover a defined contingency case. A ratio above one does not protect against common-cause defects, wrong configuration or inaccessible storage. Use the result only to justify: Count only spares that are serviceable, reachable and configuration-compatible in the acceptance matrix.

19 — Sensitivity or limit case

One quarantined spare reduces coverage from 1.5 to 1.0 in this case.

20 — Practice

Guided exercise — Critical spare coverage ratio. Four serviceable spares are available; the case needs three. Find coverage.

Guided correction — Critical spare coverage ratio

1. $C = 4/3 \approx 1.33$.
2. Then verify serial/configuration compatibility and retrieval time.

Autonomous exercise — Critical spare coverage ratio. Build a second case from “One quarantined spare reduces coverage from 1.5 to 1.0 in this case.”. Re-evaluate $C_spare = N_serviceable / N_required_case$. Name the changed input. Decide whether “Count only spares that are serviceable, reachable and configuration-compatible in the acceptance matrix.” still follows.

Autonomous correction — Critical spare coverage ratio

For *Critical spare coverage ratio*, state the altered case. Preserve $units/units = dimensionless$. Match the direction in “One quarantined spare reduces coverage from 1.5 to 1.0 in this case.”. Respect “A ratio above one does not protect against common-cause defects, wrong configuration or inaccessible storage.”. Finish by retaining or revising: Count only spares that are serviceable, reachable and configuration-compatible in the acceptance matrix.

21 — Mission decision

Count only spares that are serviceable, reachable and configuration-compatible in the acceptance matrix.

Command autonomy window during communications loss

$t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$

1 — Concrete question

For **Command autonomy window during communications loss**, how does $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$ inform bounding how long the local crew can continue safely before one autonomy constraint expires and the operational choice “Pre-authorize actions only inside the smallest validated autonomy horizon, with explicit HOLD triggers before it expires.”?

2 — Intuition without symbols

Intuition. During a communications outage, the crew can act independently only until the first protected limit expires. Procedures, resources and navigation knowledge may each allow different durations, so the shortest safe horizon governs.

3 — Quantities first

t_{proc} is procedure/authority horizon; t_{resource} protected resource horizon; t_{nav} navigation/state-estimation horizon; the minimum is the binding autonomy window.

4 — Formula

$t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$

5 — Read aloud

“ $t_{\text{auto_cmd}}$ autonomous command equals the minimum of procedure, resource and navigation horizons.”

6 — Symbols

Symbol map for Command autonomy window during communications loss. t_{proc} is procedure/authority horizon; t_{resource} protected resource horizon; t_{nav} navigation/state-estimation horizon; the minimum is the binding autonomy window.

7 — Pronunciation

Pronunciation. Say $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$. For *Command autonomy window during communications loss*, use the step-three names tied to bounding how long the local crew can continue safely before one autonomy constraint expires. Speak each *Command autonomy window during communications loss* unit with the quantity it measures.

8 — Units

$\min(\text{hours}, \text{hours}, \text{hours}) = \text{hours}$

9 — Convention

Convention. For *Command autonomy window during communications loss*, keep bounding how long the local crew can continue safely before one autonomy constraint expires on one declared boundary. Apply $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$ under that convention. The three horizons must be defined for the same mission state; communications loss can also degrade knowledge, authority and coordination progressively.

10 — Why this operation

Why this operation. $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$ answers the *Command autonomy window during communications loss* question because it represents bounding how long the local crew can continue safely before one autonomy constraint expires. In this case it yields: Navigation/state-estimation constraints limit autonomous continuation to 12 h in this scenario.

11 — Assumptions

Assumptions. Treat the *Command autonomy window during communications loss* values as one teaching case. For bounding how long the local crew can continue safely before one autonomy constraint expires, keep a single physical or operational boundary. The three horizons must be defined for the same mission state; communications loss can also degrade knowledge, authority and coordination progressively.

12 — Unit check

Unit check. Reduce $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$ for *Command autonomy window during communications loss*. The required dimension is $\min(\text{hours}, \text{hours}, \text{hours}) = \text{hours}$. A different dimension invalidates “Navigation/state-estimation constraints limit autonomous continuation to 12 h in this scenario.”.

13 — Numerical case

$t_{\text{proc}} = 18 \text{ h}$

$t_{\text{resource}} = 30 \text{ h}$

$t_{\text{nav}} = 12 \text{ h}$

$t_{\text{auto_cmd}} = \min(18, 30, 12) = 12 \text{ h}$

14 — Why each operation

Why each operation. For *Command autonomy window during communications loss*, substitute $t_{\text{proc}} = 18 \text{ h}$; $t_{\text{resource}} = 30 \text{ h}$; $t_{\text{nav}} = 12 \text{ h}$; $t_{\text{auto_cmd}} = \min(18, 30, 12) = 12 \text{ h}$ into $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$. Then verify the independent statement “12 h is less than or equal to all three component horizons”.

15 — Algebra check

Algebra check. Reverse $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$ for *Command autonomy window during communications loss* using “12 h is less than or equal to all three component horizons”. The recovered input should follow “Extending a non-limiting 30 h resource horizon does not change the 12 h minimum.”. If not, recheck units and boundaries.

16 — Mental estimate

Mental estimate. Round the dominant inputs for *Command autonomy window during communications loss*. Compare that rough scale with “Navigation/state-estimation constraints limit autonomous continuation to 12 h in this scenario.”. If they diverge sharply, inspect $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$ for units, signs or boundaries.

17 — Interpretation

Interpretation. For *Command autonomy window during communications loss*, Navigation/state-estimation constraints limit autonomous continuation to 12 h in this scenario. Operationally: Pre-authorize actions only inside the smallest validated autonomy horizon, with explicit HOLD triggers before it expires. The interpretation remains limited by “The three horizons must be defined for the same mission state; communications loss can also degrade knowledge, authority and coordination progressively.”.

18 — What it does not prove

What it does not prove. *Command autonomy window during communications loss* cannot support claims outside bounding how long the local crew can continue safely before one autonomy constraint expires. The three horizons must be defined for the same mission state; communications loss can also degrade knowledge, authority and coordination progressively. Use the result only to justify: Pre-authorize actions only inside the smallest validated autonomy horizon, with explicit HOLD triggers before it expires.

19 — Sensitivity or limit case

Extending a non-limiting 30 h resource horizon does not change the 12 h minimum.

20 — Practice

Guided exercise — Command autonomy window during communications loss.

Procedure horizon=10 h, resource horizon=14 h, navigation horizon=16 h. Find autonomy window.

Guided correction — Command autonomy window during communications loss

1. $\min=10$ h.
2. Procedure/authority is the first binding constraint.

Autonomous exercise — Command autonomy window during communications loss.

Build a second case from “Extending a non-limiting 30 h resource horizon does not change the 12 h minimum.”. Re-evaluate $t_{\text{auto_cmd}} = \min(t_{\text{proc}}, t_{\text{resource}}, t_{\text{nav}})$. Name the changed input. Decide whether “Pre-authorize actions only inside the smallest validated autonomy horizon, with explicit HOLD triggers before it expires.” still follows.

Autonomous correction — Command autonomy window during communications loss

For *Command autonomy window during communications loss*, state the altered case. Preserve $\min(\text{hours}, \text{hours}, \text{hours}) = \text{hours}$. Match the direction in “Extending a non-limiting 30 h resource horizon does not change the 12 h minimum.”. Respect “The three horizons must be defined for the same mission state; communications loss can also degrade knowledge, authority and coordination progressively.”. Finish by retaining or revising: Pre-authorize actions only inside the smallest validated autonomy horizon, with explicit HOLD triggers before it expires.

21 — Mission decision

Pre-authorize actions only inside the smallest validated autonomy horizon, with explicit HOLD triggers before it expires.

Mission-integration quantitative closure — eight competencies added for FR↔EN parity

This extension closes the eight quantitative families still present in the French final-mission course but absent as complete English mini-lessons. The purpose is not to inflate formula counts: each relationship is tied to an explicit mission decision, a dimensional check and two corrected exercises.

Delivered usable mass — distinguish gross payload from available payload

$$M_{\text{del}} = M_{\text{gross}} \times f_{\text{available}}$$

1 — Concrete question

What mission decision becomes clearer when mass that remains genuinely available after applying a declared availability factor is calculated explicitly?

2 — Intuition without symbols

Intuition. A manifest can list more mass than the mission can truly count on. Applying an explicit availability fraction separates the attractive gross number from the mass that is actually usable for planning.

3 — Quantities first

M_{gross} is nominal gross mass; $f_{\text{available}}$ is the usable fraction; M_{del} is delivered usable mass.

4 — Formula

$$M_{\text{del}} = M_{\text{gross}} \times f_{\text{available}}$$

5 — Read aloud

“M delivered equals M gross times f available.”

6 — Symbols

M_{gross} is nominal gross mass; $f_{\text{available}}$ is the usable fraction; M_{del} is delivered usable mass.

7 — Pronunciation

Pronunciation. Say $M_{\text{del}} = M_{\text{gross}} \times f_{\text{available}}$. For *Delivered usable mass — distinguish gross payload from available payload*, use these quantity names: M_{gross} is nominal gross mass; $f_{\text{available}}$ is the usable fraction; M_{del} is delivered usable mass.

8 — Units

$$\text{kg} \times \text{dimensionless} = \text{kg}$$

9 — Convention

Convention. Keep *Delivered usable mass — distinguish gross payload from available payload* on one mission boundary. Apply $M_{del} = M_{gross} \times f_{available}$ with the meanings “ M_{gross} is nominal gross mass; $f_{available}$ is the usable fraction; M_{del} is delivered usable mass.”.

10 — Why this operation

Why this operation. In *Delivered usable mass — distinguish gross payload from available payload*, $M_{del} = M_{gross} \times f_{available}$ follows the accounting defined by “ M_{gross} is nominal gross mass; $f_{available}$ is the usable fraction; M_{del} is delivered usable mass.”.

11 — Assumptions

Assumptions. The gross manifest and the availability fraction must describe the same payload boundary and the same review state. Treat the availability factor as one declared planning adjustment; do not let it hide a second mass margin, packaging loss or a separate landing reserve.

12 — Unit check

Unit check. For *Delivered usable mass — distinguish gross payload from available payload*, $M_{del} = M_{gross} \times f_{available}$ must reduce to $\text{kg} \times \text{dimensionless} = \text{kg}$. Reject any other dimension.

13 — Numerical case

$$M_{gross} = 1,800 \text{ kg}$$

$$f_{available} = 0.92$$

$$M_{del} = 1,800 \times 0.92 = 1,656 \text{ kg}$$

14 — Why each operation

Why each operation. Follow $M_{del} = M_{gross} \times f_{available}$ for *Delivered usable mass — distinguish gross payload from available payload*. Keep the units in $\text{kg} \times \text{dimensionless} = \text{kg}$ attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $M_{del} = M_{gross} \times f_{available}$ for *Delivered usable mass — distinguish gross payload from available payload*. The direction should agree with “Reducing availability from 0.92 to 0.80 lowers usable mass even though gross mass is unchanged.”.

16 — Mental estimate

Mental estimate. Round the dominant *Delivered usable mass — distinguish gross payload from available payload* inputs. Compare that rough answer with $M_{del} = M_{gross} \times f_{available}$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

The planning mass available to the mission is 1,656 kg. Operational consequence: Budget downstream systems against usable delivered mass, not against gross manifest mass.

18 — What it does not prove

What it does not prove. This calculation is an accounting adjustment, not a transport-performance model. It does not predict launch losses, landing dispersion, packaging constraints, damage on arrival or whether the delivered hardware is usable in the required configuration. Its legitimate use is to keep downstream allocations tied to the mass that planning actually makes available.

19 — Sensitivity or limit case

Reducing availability from 0.92 to 0.80 lowers usable mass even though gross mass is unchanged.

20 — Practice

Guided exercise. A manifest lists 2,200 kg gross with an available fraction of 0.88. Find the usable delivered mass.

Guided correction — Delivered usable mass — distinguish gross payload from available payload

1. $M_{del} = 2,200 \times 0.88 = 1,936$ kg.
2. The 264 kg difference is unavailable under the stated planning convention.

Autonomous exercise. A later review keeps gross mass at 2,200 kg but reduces the usable fraction to 0.80. Recalculate usable mass and state the planning consequence.

Autonomous correction — Delivered usable mass — distinguish gross payload from available payload

1. $M_{del} = 2,200 \times 0.80 = 1,760$ kg.
2. Usable mass falls by 176 kg relative to the guided case, so downstream allocations must be reopened.

21 — Mission decision

Budget downstream systems against usable delivered mass, not against gross manifest mass.

Total mission mass with explicit margin

$$M_{\text{total}} = \sum M_i \times (1 + m_{\text{mass}})$$

1 — Concrete question

What mission decision becomes clearer when rolling subsystem masses into a mission total while keeping a visible design margin is calculated explicitly?

2 — Intuition without symbols

Intuition. Adding subsystem masses gives the current design, not the protected design. A visible top-level margin reserves capacity for uncertainty so later growth does not immediately break the architecture.

3 — Quantities first

M_i is each subsystem mass; $\sum$ adds those masses; m_{mass} is the fractional mass margin; M_{total} is the margined total.

4 — Formula

$$M_{\text{total}} = \sum M_i \times (1 + m_{\text{mass}})$$

5 — Read aloud

"M total equals the sum of M i times one plus m mass."

6 — Symbols

M_i is each subsystem mass; $\sum$ adds those masses; m_{mass} is the fractional mass margin; M_{total} is the margined total.

7 — Pronunciation

Pronunciation. Say $M_{\text{total}} = \sum M_i \times (1 + m_{\text{mass}})$. For *Total mission mass with explicit margin*, use these quantity names: M_i is each subsystem mass; $\sum$ adds those masses; m_{mass} is the fractional mass margin; M_{total} is the margined total.

8 — Units

$$\text{kg} \times \text{dimensionless} = \text{kg}$$

9 — Convention

Convention. Keep *Total mission mass with explicit margin* on one mission boundary. Apply $M_{\text{total}} = \sum M_i \times (1 + m_{\text{mass}})$ with the meanings " M_i is each subsystem mass; $\sum$ adds those masses; m_{mass} is the fractional mass margin; M_{total} is the margined total."

10 — Why this operation

Why this operation. In *Total mission mass with explicit margin*, $M_{total} = \sum M_i \times (1 + m_{mass})$ follows the accounting defined by “ M_i is each subsystem mass; $\sum$ adds those masses; m_{mass} is the fractional mass margin; M_{total} is the margined total.”.

11 — Assumptions

Assumptions. Apply the top-level margin only to subsystem masses that have been declared on a compatible basis. If a subsystem already contains the same reserve, remove that reserve before the top-level calculation or the design will double-count protection.

12 — Unit check

Unit check. For *Total mission mass with explicit margin*, $M_{total} = \sum M_i \times (1 + m_{mass})$ must reduce to $\text{kg} \times \text{dimensionless} = \text{kg}$. Reject any other dimension.

13 — Numerical case

$$\sum M_i = 10,500 \text{ kg}$$

$$m_{mass} = 0.15$$

$$M_{total} = 10,500 \times 1.15 = 12,075 \text{ kg}$$

14 — Why each operation

Why each operation. Follow $M_{total} = \sum M_i \times (1 + m_{mass})$ for *Total mission mass with explicit margin*. Keep the units in $\text{kg} \times \text{dimensionless} = \text{kg}$ attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $M_{total} = \sum M_i \times (1 + m_{mass})$ for *Total mission mass with explicit margin*. The direction should agree with “A larger margin increases total mass linearly when the unmargined subsystem sum is fixed.”.

16 — Mental estimate

Mental estimate. Round the dominant *Total mission mass with explicit margin* inputs. Compare that rough answer with $M_{total} = \sum M_i \times (1 + m_{mass})$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

The margined mission total is 12,075 kg. Operational consequence: Expose margin separately so launch and landing trades can distinguish hardware growth from protected reserve.

18 — What it does not prove

What it does not prove. A margined mass total does not establish launch feasibility, volume closure, centre-of-mass compliance, entry or landing performance. It only makes the mass reserve visible. Vehicle and trajectory limits still need their own gates before the mission can be accepted.

19 — Sensitivity or limit case

A larger margin increases total mass linearly when the unmargined subsystem sum is fixed.

20 — Practice

Guided exercise. Subsystems sum to 8,400 kg and the review board requires a 20% mass margin. Find total margined mass.

Guided correction — Total mission mass with explicit margin

1. $M_{\text{total}} = 8,400 \times 1.20 = 10,080$ kg.
2. The margin contributes 1,680 kg above the unmargined sum.

Autonomous exercise. A redesign produces 9,100 kg of unmargined subsystem mass with a 10% top-level margin. Compute the new total and compare it with the guided case.

Autonomous correction — Total mission mass with explicit margin

1. $M_{\text{total}} = 9,100 \times 1.10 = 10,010$ kg.
2. Despite the higher base mass, the smaller margin yields a slightly lower total than 10,080 kg.

21 — Mission decision

Expose margin separately so launch and landing trades can distinguish hardware growth from protected reserve.

Daily energy budget — add power use over operating time

$$E_{\text{day}} = \sum (P_i \times t_i)$$

1 — Concrete question

What mission decision becomes clearer when daily electrical-energy demand across loads with different powers and duty times is calculated explicitly?

2 — Intuition without symbols

Intuition. Power tells how fast energy is being used; duration tells how long that draw continues. Multiplying each load by its operating time and then adding the contributions gives the energy the day must supply.

3 — Quantities first

P_i is the power of load i ; t_i is its operating duration; E_{day} is the summed daily energy.

4 — Formula

$$E_{\text{day}} = \sum (P_i \times t_i)$$

5 — Read aloud

“ E_{day} equals the sum of P_i times t_i .”

6 — Symbols

P_i is the power of load i ; t_i is its operating duration; E_{day} is the summed daily energy.

7 — Pronunciation

Pronunciation. Say $E_{\text{day}} = \sum (P_i \times t_i)$. For *Daily energy budget — add power use over operating time*, use these quantity names: P_i is the power of load i ; t_i is its operating duration; E_{day} is the summed daily energy.

8 — Units

$$\text{kW} \times \text{h} = \text{kWh}$$

9 — Convention

Convention. Keep *Daily energy budget — add power use over operating time* on one mission boundary. Apply $E_{\text{day}} = \sum (P_i \times t_i)$ with the meanings “ P_i is the power of load i ; t_i is its operating duration; E_{day} is the summed daily energy.”.

10 — Why this operation

Why this operation. In *Daily energy budget — add power use over operating time*, $E_{\text{day}} = \sum (P_i \times t_i)$ follows the accounting defined by “ P_i is the power of load i ; t_i is its operating duration; E_{day} is the summed daily energy.”.

11 — Assumptions

Assumptions. Every power-duration pair must refer to the same accounting day and to a clearly defined operating state. Conversion losses, standby loads and cycling penalties belong in the ledger only when the stated power values actually include them.

12 — Unit check

Unit check. For *Daily energy budget — add power use over operating time*, $E_{\text{day}} = \sum(P_i \times t_i)$ must reduce to $\text{kW} \times \text{h} = \text{kWh}$. Reject any other dimension.

13 — Numerical case

Load A: $2 \text{ kW} \times 10 \text{ h} = 20 \text{ kWh}$

Load B: $4 \text{ kW} \times 6 \text{ h} = 24 \text{ kWh}$

Load C: $1 \text{ kW} \times 24 \text{ h} = 24 \text{ kWh}$

$E_{\text{day}} = 68 \text{ kWh}$

14 — Why each operation

Why each operation. Follow $E_{\text{day}} = \sum(P_i \times t_i)$ for *Daily energy budget — add power use over operating time*. Keep the units in $\text{kW} \times \text{h} = \text{kWh}$ attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $E_{\text{day}} = \sum(P_i \times t_i)$ for *Daily energy budget — add power use over operating time*. The direction should agree with “Extending the operating time of a high-power load increases the daily budget faster than extending a low-power load by the same number of hours.”.

16 — Mental estimate

Mental estimate. Round the dominant *Daily energy budget — add power use over operating time* inputs. Compare that rough answer with $E_{\text{day}} = \sum(P_i \times t_i)$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

The three-load daily energy demand is 68 kWh. Operational consequence: Size generation and storage from an auditable energy ledger, then check peak power separately.

18 — What it does not prove

What it does not prove. Summed daily energy cannot size peak generation, switchgear or an inverter because coincident loads and transients disappear inside the daily total. It is the energy ledger for generation and storage trades; peak-power and short-duration surge checks remain separate acceptance tests.

19 — Sensitivity or limit case

Extending the operating time of a high-power load increases the daily budget faster than extending a low-power load by the same number of hours.

20 — Practice

Guided exercise. One load uses 3 kW for 8 h and another uses 1.5 kW for 16 h. Find daily energy.

Guided correction — Daily energy budget — add power use over operating time

1. First load: $3 \times 8 = 24$ kWh.
2. Second load: $1.5 \times 16 = 24$ kWh.
3. Total = 48 kWh.

Autonomous exercise. A revised day uses 5 kW for 4 h, 2 kW for 10 h and 0.5 kW for 24 h. Calculate total energy and identify the largest contributor.

Autonomous correction — Daily energy budget — add power use over operating time

1. Energy contributions are 20, 20 and 12 kWh.
2. Total daily energy is 52 kWh; the first two loads tie for the largest contribution.

21 — Mission decision

Size generation and storage from an auditable energy ledger, then check peak power separately.

Expected failures from mission duration and MTBF

$$N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$$

1 — Concrete question

What mission decision becomes clearer when first-order spare planning from exposure time and mean time between failures is calculated explicitly?

2 — Intuition without symbols

Intuition. A component exposed for much longer than its typical interval between failures is expected to fail more often. Dividing exposure time by that interval gives a first-order count for planning, not a timetable of events.

3 — Quantities first

t_{mission} is exposure duration; MTBF is mean time between failures under the assumed model; N_{fail} is the expected failure count.

4 — Formula

$$N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$$

5 — Read aloud

"N fail equals t mission divided by M T B F."

6 — Symbols

t_mission is exposure duration; MTBF is mean time between failures under the assumed model; N_fail is the expected failure count.

7 — Pronunciation

Pronunciation. Say $N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$. For *Expected failures from mission duration and MTBF*, use these quantity names: t_mission is exposure duration; MTBF is mean time between failures under the assumed model; N_fail is the expected failure count.

8 — Units

h / h = dimensionless expected count

9 — Convention

Convention. Keep *Expected failures from mission duration and MTBF* on one mission boundary. Apply $N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$ with the meanings "t_mission is exposure duration; MTBF is mean time between failures under the assumed model; N_fail is the expected failure count."

10 — Why this operation

Why this operation. In *Expected failures from mission duration and MTBF*, $N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$ follows the accounting defined by "t_mission is exposure duration; MTBF is mean time between failures under the assumed model; N_fail is the expected failure count."

11 — Assumptions

Assumptions. The MTBF must be applicable to the stated operating environment and exposure interval, and this first-order estimate assumes a regime in which a constant average failure rate is a defensible approximation. Mixing unlike populations or life phases makes the quotient misleading.

12 — Unit check

Unit check. For *Expected failures from mission duration and MTBF*, $N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$ must reduce to h / h = dimensionless expected count. Reject any other dimension.

13 — Numerical case

$t_{\text{mission}} = 7,200 \text{ h}$

$\text{MTBF} = 2,400 \text{ h}$

$N_{\text{fail}} = 7,200 / 2,400 = 3.0$

14 — Why each operation

Why each operation. Follow $N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$ for *Expected failures from mission duration and MTBF*. Keep the units in $\text{h} / \text{h} = \text{dimensionless}$ expected count attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$ for *Expected failures from mission duration and MTBF*. The direction should agree with “A longer mission or shorter MTBF increases expected failures.”.

16 — Mental estimate

Mental estimate. Round the dominant *Expected failures from mission duration and MTBF* inputs. Compare that rough answer with $N_{\text{fail}} = t_{\text{mission}} / \text{MTBF}$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

The simple expectation is three failures over the stated exposure. Operational consequence: Use expected failures to screen spare demand, then validate with a reliability model appropriate to the actual failure regime.

18 — What it does not prove

What it does not prove. An expected count is not a timetable of failures and cannot represent infant mortality, wear-out, common-cause events or repair interactions. Use it only to screen spare demand and workload before applying a reliability model that matches the actual failure process.

19 — Sensitivity or limit case

A longer mission or shorter MTBF increases expected failures.

20 — Practice

Guided exercise. A component is exposed for 5,000 h with an MTBF of 2,000 h. Find the simple expected failure count.

Guided correction — Expected failures from mission duration and MTBF

1. $N_{\text{fail}} = 5,000 / 2,000 = 2.5$.
2. The fractional result is an expectation across repeated cases, not half of a physical failure.

Autonomous exercise. Another component is exposed for 8,000 h with an MTBF of 3,200 h. Compute the expectation and state why the result is not a prediction of exact failure times.

Autonomous correction — Expected failures from mission duration and MTBF

1. $N_{\text{fail}} = 8,000 / 3,200 = 2.5$.
2. The value describes the average count under the assumed constant-rate model; actual events can occur earlier, later or not at all.

21 — Mission decision

Use expected failures to screen spare demand, then validate with a reliability model appropriate to the actual failure regime.

Capacity margin — measure protected capacity above need

$$m_{\text{cap}} = (C_{\text{available}} - C_{\text{required}}) / C_{\text{required}}$$

1 — Concrete question

What mission decision becomes clearer when expressing how much capacity remains above a defined requirement is calculated explicitly?

2 — Intuition without symbols

Intuition. Extra capacity is meaningful only relative to what the mission needs. Comparing the excess with the requirement shows whether the reserve is generous or thin on the scale of the obligation.

3 — Quantities first

$C_{\text{available}}$ is available capacity; C_{required} is required capacity; m_{cap} is relative capacity margin.

4 — Formula

$$m_{\text{cap}} = (C_{\text{available}} - C_{\text{required}}) / C_{\text{required}}$$

5 — Read aloud

"m capacity equals available capacity minus required capacity, divided by required capacity."

6 — Symbols

C_available is available capacity; C_required is required capacity; m_cap is relative capacity margin.

7 — Pronunciation

Pronunciation. Say $m_cap = (C_available - C_required) / C_required$. For *Capacity margin — measure protected capacity above need*, use these quantity names: C_available is available capacity; C_required is required capacity; m_cap is relative capacity margin.

8 — Units

same unit / same unit = dimensionless

9 — Convention

Convention. Keep *Capacity margin — measure protected capacity above need* on one mission boundary. Apply $m_cap = (C_available - C_required) / C_required$ with the meanings "C_available is available capacity; C_required is required capacity; m_cap is relative capacity margin."

10 — Why this operation

Why this operation. In *Capacity margin — measure protected capacity above need*, $m_cap = (C_available - C_required) / C_required$ follows the accounting defined by "C_available is available capacity; C_required is required capacity; m_cap is relative capacity margin."

11 — Assumptions

Assumptions. Available and required capacity must measure the same function under the same environmental, duty-cycle and rating conditions. A nominal catalogue rating cannot be compared directly with a derated mission requirement unless the rating basis has first been reconciled.

12 — Unit check

Unit check. For *Capacity margin — measure protected capacity above need*, $m_cap = (C_available - C_required) / C_required$ must reduce to same unit / same unit = dimensionless. Reject any other dimension.

13 — Numerical case

$C_{\text{available}} = 125 \text{ units}$

$C_{\text{required}} = 100 \text{ units}$

$m_{\text{cap}} = (125 - 100) / 100 = 0.25 = 25\%$

14 — Why each operation

Why each operation. Follow $m_{\text{cap}} = (C_{\text{available}} - C_{\text{required}}) / C_{\text{required}}$ for *Capacity margin — measure protected capacity above need*. Keep the units in same unit / same unit = dimensionless attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $m_{\text{cap}} = (C_{\text{available}} - C_{\text{required}}) / C_{\text{required}}$ for *Capacity margin — measure protected capacity above need*. The direction should agree with “When requirement rises while available capacity stays fixed, relative margin shrinks.”.

16 — Mental estimate

Mental estimate. Round the dominant *Capacity margin — measure protected capacity above need* inputs. Compare that rough answer with $m_{\text{cap}} = (C_{\text{available}} - C_{\text{required}}) / C_{\text{required}}$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

Available capacity is 25% above the stated requirement. Operational consequence: Track margin against the requirement actually used for certification, not against an easier nominal case.

18 — What it does not prove

What it does not prove. A positive capacity margin does not demonstrate reliability, dynamic response, product quality or contingency recovery. It only quantifies headroom against one stated requirement, so certification still needs the failure, transient and degraded-mode cases that can consume that headroom.

19 — Sensitivity or limit case

When requirement rises while available capacity stays fixed, relative margin shrinks.

20 — Practice

Guided exercise. A subsystem can provide 96 units against an 80-unit requirement. Find relative margin.

Guided correction — Capacity margin — measure protected capacity above need

1. $m_{cap} = (96 - 80) / 80 = 0.20 = 20\%$.
2. The system carries 16 units of absolute excess capacity.

Autonomous exercise. A redesigned subsystem provides 138 units against a 120-unit requirement. Find margin and compare it with the guided case.

Autonomous correction — Capacity margin — measure protected capacity above need

1. $m_{cap} = (138 - 120) / 120 = 0.15 = 15\%$.
2. The absolute excess is larger than before, but the relative protected margin is smaller.

21 — Mission decision

Track margin against the requirement actually used for certification, not against an easier nominal case.

Series-system reliability — multiply independent success probabilities

$$R_{sys} = \prod R_i$$

1 — Concrete question

What mission decision becomes clearer when estimating success probability for a simplified series chain in which every element must succeed is calculated explicitly?

2 — Intuition without symbols

Intuition. A pure series chain succeeds only if every required element succeeds. Even highly reliable components therefore combine into a system reliability that is lower than each perfect-looking part would suggest alone.

3 — Quantities first

R_i is reliability of element i over the stated interval; $\prod$ multiplies all element reliabilities; R_{sys} is series reliability.

4 — Formula

$$R_{sys} = \prod R_i$$

5 — Read aloud

"R system equals the product of R i."

6 — Symbols

R_i is reliability of element i over the stated interval; Π multiplies all element reliabilities; R_{sys} is series reliability.

7 — Pronunciation

Pronunciation. Say R_{sys} = Π R_i. For *Series-system reliability — multiply independent success probabilities*, use these quantity names: R_i is reliability of element i over the stated interval; Π multiplies all element reliabilities; R_{sys} is series reliability.

8 — Units

dimensionless × dimensionless = dimensionless

9 — Convention

Convention. Keep *Series-system reliability — multiply independent success probabilities* on one mission boundary. Apply R_{sys} = Π R_i with the meanings "R_i is reliability of element i over the stated interval; Π multiplies all element reliabilities; R_{sys} is series reliability."

10 — Why this operation

Why this operation. In *Series-system reliability — multiply independent success probabilities*, R_{sys} = Π R_i follows the accounting defined by "R_i is reliability of element i over the stated interval; Π multiplies all element reliabilities; R_{sys} is series reliability."

11 — Assumptions

Assumptions. Each probability must refer to the same mission interval, every listed element must be required for success, and the multiplication assumes the element outcomes are independent for this simplified model. Repairs, standby redundancy and common-cause failures violate that model.

12 — Unit check

Unit check. For *Series-system reliability — multiply independent success probabilities*, R_{sys} = Π R_i must reduce to dimensionless × dimensionless = dimensionless. Reject any other dimension.

13 — Numerical case

$$R_1 = 0.98$$

$$R_2 = 0.97$$

$$R_3 = 0.99$$

$$R_{\text{sys}} = 0.98 \times 0.97 \times 0.99 \approx 0.941 = 94.1\%$$

14 — Why each operation

Why each operation. Follow $R_{\text{sys}} = \prod R_i$ for *Series-system reliability — multiply independent success probabilities*. Keep the units in dimensionless $\times$ dimensionless = dimensionless attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $R_{\text{sys}} = \prod R_i$ for *Series-system reliability — multiply independent success probabilities*. The direction should agree with “Adding another element with reliability below one decreases pure series reliability.”.

16 — Mental estimate

Mental estimate. Round the dominant *Series-system reliability — multiply independent success probabilities* inputs. Compare that rough answer with $R_{\text{sys}} = \prod R_i$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

The simplified three-element series reliability is about 94.1%. Operational consequence: Use the product only as a transparent first-order series model and escalate to fault-tree or reliability-block analysis when dependencies matter.

18 — What it does not prove

What it does not prove. The series product is not an availability model and cannot represent redundancy, repair, shared dependencies or cascading faults. It is a transparent first-order screen; architectures with coupled failures require a fault tree, reliability block diagram or another dependency-aware model.

19 — Sensitivity or limit case

Adding another element with reliability below one decreases pure series reliability.

20 — Practice

Guided exercise. Two independent series elements have reliabilities 0.95 and 0.98. Find system reliability.

Guided correction — Series-system reliability — multiply independent success probabilities

1. $R_{\text{sys}} = 0.95 \times 0.98 = 0.931$.
2. The series system reliability is 93.1%.

Autonomous exercise. Three series elements have reliabilities 0.99, 0.99 and 0.98. Compute system reliability and identify why a high component value does not guarantee an equally high system value.

Autonomous correction — Series-system reliability — multiply independent success probabilities

1. $R_{\text{sys}} = 0.99 \times 0.99 \times 0.98 \approx 0.9605 = 96.05\%$.
2. Every required element introduces another chance of failure, so the series product falls below the best individual values.

21 — Mission decision

Use the product only as a transparent first-order series model and escalate to fault-tree or reliability-block analysis when dependencies matter.

Crew water inventory — convert daily demand into campaign mass

$$m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_{\text{s}}$$

1 — Concrete question

What mission decision becomes clearer when planning gross water demand before recycling and make-up flows are applied is calculated explicitly?

2 — Intuition without symbols

Intuition. Water demand grows with three simple things: how much one person needs, how many people must be supported and how long support must last. Multiplying those dimensions gives the gross campaign requirement before recycling is credited.

3 — Quantities first

q_{water} is water demand per crewmember per sol or day; N_{crew} is crew size; t_{s} is covered duration; m_{water} is gross water demand.

4 — Formula

$$m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_{\text{s}}$$

5 — Read aloud

“m_{water} equals q_{water} times N_{crew} times t_s.”

6 — Symbols

q_{water} is water demand per crewmember per sol or day; N_{crew} is crew size; t_s is covered duration; m_{water} is gross water demand.

7 — Pronunciation

Pronunciation. Say $m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_s$. For *Crew water inventory — convert daily demand into campaign mass*, use these quantity names: q_{water} is water demand per crewmember per sol or day; N_{crew} is crew size; t_s is covered duration; m_{water} is gross water demand.

8 — Units

$\text{kg}/(\text{person} \cdot \text{sol}) \times \text{person} \times \text{sol} = \text{kg}$

9 — Convention

Convention. Keep *Crew water inventory — convert daily demand into campaign mass* on one mission boundary. Apply $m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_s$ with the meanings “q_{water} is water demand per crewmember per sol or day; N_{crew} is crew size; t_s is covered duration; m_{water} is gross water demand.”.

10 — Why this operation

Why this operation. In *Crew water inventory — convert daily demand into campaign mass*, $m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_s$ follows the accounting defined by “q_{water} is water demand per crewmember per sol or day; N_{crew} is crew size; t_s is covered duration; m_{water} is gross water demand.”.

11 — Assumptions

Assumptions. The per-person demand rate, crew count and covered duration must use the same campaign boundary. This is deliberately a gross-demand calculation: recycling credit, inaccessible inventory, leakage and protected reserve are introduced only in the subsequent closed-loop budget.

12 — Unit check

Unit check. For *Crew water inventory — convert daily demand into campaign mass*, $m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_s$ must reduce to $\text{kg}/(\text{person} \cdot \text{sol}) \times \text{person} \times \text{sol} = \text{kg}$. Reject any other dimension.

13 — Numerical case

$$q_{\text{water}} = 3.2 \text{ kg}/(\text{person} \cdot \text{sol})$$

$$N_{\text{crew}} = 6$$

$$t_{\text{s}} = 30 \text{ sol}$$

$$m_{\text{water}} = 3.2 \times 6 \times 30 = 576 \text{ kg}$$

14 — Why each operation

Why each operation. Follow $m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_{\text{s}}$ for *Crew water inventory — convert daily demand into campaign mass*. Keep the units in $\text{kg}/(\text{person} \cdot \text{sol}) \times \text{person} \times \text{sol} = \text{kg}$ attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_{\text{s}}$ for *Crew water inventory — convert daily demand into campaign mass*. The direction should agree with “Every additional crew member or covered sol increases gross demand linearly when the per-person rate is fixed.”.

16 — Mental estimate

Mental estimate. Round the dominant *Crew water inventory — convert daily demand into campaign mass* inputs. Compare that rough answer with $m_{\text{water}} = q_{\text{water}} \times N_{\text{crew}} \times t_{\text{s}}$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

Gross water demand for the example campaign is 576 kg. Operational consequence: Close gross demand first, then layer recovery, loss and contingency assumptions transparently.

18 — What it does not prove

What it does not prove. Gross water demand does not equal imported or locally produced make-up water. Recovery efficiency, losses, storage accessibility, contingency reserve and process-water needs can all change the amount that must actually be supplied to the habitat.

19 — Sensitivity or limit case

Every additional crew member or covered sol increases gross demand linearly when the per-person rate is fixed.

20 — Practice

Guided exercise. Four crewmembers require 2.8 kg per person per sol for 20 sols. Find gross water demand.

Guided correction — Crew water inventory — convert daily demand into campaign mass

1. $m_{\text{water}} = 2.8 \times 4 \times 20 = 224 \text{ kg}$.
2. This is demand before crediting recycling.

Autonomous exercise. Five crewmembers are planned at 3.0 kg per person per sol for 40 sols. Compute gross demand and state what additional information is needed to convert it into make-up water.

Autonomous correction — Crew water inventory — convert daily demand into campaign mass

1. $m_{\text{water}} = 3.0 \times 5 \times 40 = 600 \text{ kg}$.
2. A recovery fraction, losses and protected reserve are needed before translating gross demand into imported or locally produced make-up water.

21 — Mission decision

Close gross demand first, then layer recovery, loss and contingency assumptions transparently.

Energy per Martian sol — convert average power into sol energy

$$E_{\text{sol}} = P_{\text{avg}} \times t_{\text{sol}}$$

1 — Concrete question

What mission decision becomes clearer when converting an average continuous power level into energy over one Martian sol is calculated explicitly?

2 — Intuition without symbols

Intuition. A steady average power draw accumulates energy throughout the Martian day. Multiplying that average level by the length of a sol converts an instantaneous rate into the energy budget that generation and storage must deliver.

3 — Quantities first

P_{avg} is average power over the interval; t_{sol} is sol duration; E_{sol} is energy over one sol.

4 — Formula

$$E_{\text{sol}} = P_{\text{avg}} \times t_{\text{sol}}$$

5 — Read aloud

"E_sol equals P_avg times t_sol."

6 — Symbols

P_avg is average power over the interval; t_sol is sol duration; E_sol is energy over one sol.

7 — Pronunciation

Pronunciation. Say $E_{sol} = P_{avg} \times t_{sol}$. For *Energy per Martian sol — convert average power into sol energy*, use these quantity names: P_avg is average power over the interval; t_sol is sol duration; E_sol is energy over one sol.

8 — Units

$\text{kW} \times \text{h} = \text{kWh}$

9 — Convention

Convention. Keep *Energy per Martian sol — convert average power into sol energy* on one mission boundary. Apply $E_{sol} = P_{avg} \times t_{sol}$ with the meanings "P_avg is average power over the interval; t_sol is sol duration; E_sol is energy over one sol."

10 — Why this operation

Why this operation. In *Energy per Martian sol — convert average power into sol energy*, $E_{sol} = P_{avg} \times t_{sol}$ follows the accounting defined by "P_avg is average power over the interval; t_sol is sol duration; E_sol is energy over one sol."

11 — Assumptions

Assumptions. The average power must genuinely represent the full Martian-sol interval used in the calculation. The example uses a sol duration of 24.6597 hours; any shorter operating window or duty-cycled load must be represented by its own time history rather than hidden inside that average.

12 — Unit check

Unit check. For *Energy per Martian sol — convert average power into sol energy*, $E_{sol} = P_{avg} \times t_{sol}$ must reduce to $\text{kW} \times \text{h} = \text{kWh}$. Reject any other dimension.

13 — Numerical case

$P_{avg} = 5.0 \text{ kW}$

$t_{sol} = 24.6597 \text{ h}$

$E_{sol} = 5.0 \times 24.6597 \approx 123.3 \text{ kWh}$

14 — Why each operation

Why each operation. Follow $E_{\text{sol}} = P_{\text{avg}} \times t_{\text{sol}}$ for *Energy per Martian sol* — convert average power into sol energy. Keep the units in $\text{kW} \times \text{h} = \text{kWh}$ attached to each intermediate value.

15 — Algebra check

Algebra check. Reverse one operation in $E_{\text{sol}} = P_{\text{avg}} \times t_{\text{sol}}$ for *Energy per Martian sol* — convert average power into sol energy. The direction should agree with “At fixed sol duration, doubling average power doubles energy per sol.”.

16 — Mental estimate

Mental estimate. Round the dominant *Energy per Martian sol* — convert average power into sol energy inputs. Compare that rough answer with $E_{\text{sol}} = P_{\text{avg}} \times t_{\text{sol}}$. An order-of-magnitude disagreement needs investigation.

17 — Interpretation

A continuous 5 kW average corresponds to about 123.3 kWh per Martian sol. Operational consequence: Use sol energy for storage and generation ledgers, while retaining a separate peak-power gate.

18 — What it does not prove

What it does not prove. Energy per sol cannot reveal instantaneous peak power, ramp rate or the timing of storage charge and discharge. Use it to close the sol-scale energy ledger, then retain separate gates for peak generation, inverter capacity and transient storage behaviour.

19 — Sensitivity or limit case

At fixed sol duration, doubling average power doubles energy per sol.

20 — Practice

Guided exercise. A subsystem averages 3.5 kW across a Martian sol. Using 24.6597 h, find energy per sol.

Guided correction — Energy per Martian sol — convert average power into sol energy

1. $E_{\text{sol}} = 3.5 \times 24.6597 \approx 86.3 \text{ kWh}$.
2. The value is an energy total, not the subsystem peak power.

Autonomous exercise. A habitat averages 8.0 kW across one sol. Compute energy and state why this number alone cannot size an inverter.

Autonomous correction — Energy per Martian sol — convert average power into sol energy

1. $E_{\text{sol}} = 8.0 \times 24.6597 \approx 197.3 \text{ kWh}$.
2. Inverter sizing depends on instantaneous and coincident peak power, which the average-energy figure does not contain.

21 — Mission decision

Use sol energy for storage and generation ledgers, while retaining a separate peak-power gate.

FR↔EN quantitative competency matrix — module 16

The matrix below maps the eighteen French quantitative families to their English teaching location. Duplicate French formulations are mapped to the same underlying English competency when the physical skill is identical.

French competency family	English teaching block	R74 status
Minimum protected stock	Minimum protected stock	PASS — existing EN
Need with margin	Minimum protected stock / explicit mass margin	PASS — covered across EN stock + R74 mass margin
Delivered usable mass	Delivered usable mass	PASS — R74
Total mass with margin	Total mission mass with explicit margin	PASS — R74
Daily energy ledger	Daily energy budget	PASS — R74
Crew stock sizing	Minimum protected stock	PASS — existing EN
Delivered mass duplicate in FR	Delivered usable mass	PASS — R74
Stock autonomy	Oxygen and food reserve horizons	PASS — existing EN
Rescue geometry	Protected rescue radius	PASS — existing EN

French competency family	English teaching block	R74 status
Expected failures / MTBF	Expected failures from mission duration and MTBF	PASS — R74
Capacity margin	Capacity margin	PASS — R74
Series reliability	Series-system reliability	PASS — R74
Water inventory	Crew water inventory	PASS — R74
Oxygen inventory / horizon	Crew oxygen reserve horizon	PASS — existing EN
Food inventory / horizon	Food stock horizon	PASS — existing EN
Energy per Martian sol	Energy per Martian sol	PASS — R74
Rover protected radius	Protected rescue radius	PASS — existing EN
Qualified human-hours	Maintenance workload utilization	PASS — existing EN

Mission systems engineering qualification — requirements, interfaces, V&V and the 500-sol proof chain

The final mission module already asks the learner to close mass, energy, consumables, maintenance, rescue and reliability. R75 adds the professional discipline that makes those budgets auditable across an integrated architecture. A mission is not ready because every subsystem owner can defend a slide. It is ready when stakeholder expectations have been translated into traceable requirements, interfaces are controlled, margins are visible, verification shows that each requirement is satisfied, validation shows that the resulting system can accomplish the intended mission, and configuration control preserves the evidence after changes.

NASA's Systems Engineering Handbook distinguishes product verification from product validation and treats verification methods such as analysis, inspection, demonstration and test as planned evidence. The same handbook emphasizes requirements, interface management, technical performance measures and end-to-end integration. The teaching architecture below adapts those ideas to a Mars campaign without pretending that a classroom exercise is a flight certification.

Requirements

Turn expectations into statements that are necessary, measurable and traceable. A requirement that cannot be verified is not ready to control design.

Interfaces

Record what crosses subsystem boundaries: mass, power, fluids, data, loads, geometry, timing, authority and failure propagation.

Verification

Establish objective evidence that the implemented product satisfies its specified requirements.

Validation

Establish that the integrated system, procedures and people can accomplish the intended mission in the relevant operational context.

1. Requirements traceability starts with the mission need and ends with evidence

A requirement should exist because a higher-level mission need demands it, not because a subsystem designer wants to freeze a preferred solution. For a Mars architecture, a stakeholder expectation such as “the crew shall retain a survivable atmosphere after a single credible cabin-isolation event” must be decomposed into measurable requirements on compartmentation, valves, sensors, control software, emergency oxygen, procedures and verification. Every child requirement should trace upward to a mission reason and downward to a verification method.

Traceability protects the project in two directions. Upward traceability exposes orphan requirements that consume mass or cost without a mission justification. Downward traceability exposes mission needs that have no implemented requirement or no planned evidence. R75 therefore treats the requirements matrix as a living engineering product, not paperwork created after the design is finished.

2. A verification matrix prevents “we tested it” from becoming meaningless

Different requirements demand different evidence. A dimensional envelope may be verified by inspection; a structural margin by analysis plus qualification testing; a crew procedure by demonstration in a representative simulator; software timing by test and analysis. The method must be chosen before the verification campaign so the project knows what evidence will close each requirement and what configuration the evidence applies to.

ID	Teaching requirement	Primary method	Closure evidence
M16-R01	Critical life-support loads shall remain powered for the protected emergency interval.	Analysis + test	Power model with worst-case loads; integrated battery/bus test at the released configuration.
M16-R02	The crew shall be able to isolate a leaking cabin volume within the response allocation.	Demonstration	Timed integrated drill with representative controls, alarms and crew protective equipment.
M16-R03	The surface mobility concept shall preserve protected return capability at the declared operating radius.	Analysis + demonstration	Energy/terrain model plus representative degraded-return field exercise.
M16-R04	Critical spare strategy shall cover the declared protected failure set.	Inspection + analysis	Configuration-controlled spare list linked to failure modes and replacement procedures.
M16-R05	The command architecture shall support the protected autonomy interval during communications loss.	Demonstration + test	Blackout thread test exercising procedures, onboard data and authority boundaries.

ID	Teaching requirement	Primary method	Closure evidence
M16-R06	Emergency consumable accounting shall remain valid after any approved configuration change.	Analysis + inspection	Recomputed budgets and signed change record showing impacts on every dependent ledger.
M16-R07	Mission data required for handover shall remain reproducible by the next crew.	Demonstration	Independent handover exercise in which a new team reconstructs configuration, open risks and protected margins.
M16-R08	Integrated mission operations shall tolerate the declared common-cause loss without violating crew-survival constraints.	End-to-end test + analysis	Integrated scenario evidence plus model-backed confirmation of untestable extremes.

3. Verification and validation answer different questions

Verification asks whether the product was built to the specified requirements. Validation asks whether the resulting product is capable of satisfying the intended mission and stakeholder expectations in its relevant environment and operational concept. The distinction matters because a perfectly verified subsystem can still participate in a mission that is operationally wrong. A habitat may meet every component specification yet impose maintenance workload that the crew cannot sustain; a rover can meet range requirements on a benign test course yet fail the intended rescue concept on representative terrain.

Module 16 therefore requires a validation argument that crosses hardware, software, procedures, crew skills, support equipment, logistics and environmental assumptions. Evidence can combine high-fidelity simulation, integrated demonstrations, analogue operations and analyses. The goal is not to recreate Mars perfectly on Earth; it is to identify which mission expectations remain at risk and use the most relevant evidence available for each.

4. Interface control is where subsystem optimism collides with reality

Many mission failures begin at boundaries. One subsystem assumes another provides a clean power bus; one software team assumes a message arrives within a latency that the communications architecture cannot guarantee; a habitat assumes a maintenance clearance that packed cargo removes; a rover assumes a charging connector can be reached while a pressure suit is worn. Interface control documents or equivalent authoritative interface products force those assumptions into explicit, versioned agreements.

For the Mars architecture, control at least mechanical geometry, structural loads, electrical voltage/current and fault behaviour, thermal rejection, fluid composition/pressure/flow, data formats and timing, communications authority, human access, contamination boundaries and emergency-isolation behaviour. Each interface should state ownership and change authority. If two teams can change the same boundary independently, the architecture does not actually control that boundary.

5. Margins are managed reserves, not decorative percentages

NASA systems engineering guidance treats margins as allowances carried in budgets and performance parameters to account for uncertainty and risk. A Mars mission should therefore distinguish current best estimate from allocation, reserve and protected minimum. Mass margin, power margin, data/storage margin, consumable reserve, schedule reserve and crew-time reserve have different physical meanings and different owners.

A recurring anti-pattern is double counting: a subsystem reports conservative demand while the system team adds another margin without understanding what uncertainty is already embedded, or two subsystems both assume access to the same central reserve. R75 requires a margin policy that records baseline value, uncertainty basis, allocation, owner, release authority and trend. A margin that falls over time without a corresponding risk decision is a warning signal, not normal project maturity.

6. Technical performance measures must show trends before a limit is crossed

A technical performance measure is useful when it reveals whether the architecture is converging. For this teaching mission, track at least total delivered mass, continuous and peak power, emergency energy horizon, heat-rejection load, oxygen/water/food protected horizons, maintenance utilization, critical spare coverage, communication-autonomy interval and return/rescue margin. The board should see the trajectory of each measure across configuration releases, not only the latest number.

Trend review changes the conversation. A power system still showing positive margin can deserve intervention if every design iteration consumes that margin faster than expected. Conversely, a low but stable margin backed by mature evidence may be less dangerous than a large nominal margin based on uncertain assumptions. The project therefore pairs the numerical value with confidence in the model and the open actions that could change it.

7. Qualification, acceptance and certification are not interchangeable labels

NASA's Systems Engineering Handbook distinguishes verification, qualification and acceptance. Qualification demonstrates that a design can meet requirements over the anticipated environmental extremes with margin. Acceptance applies a selected subset of checks to a particular flight item or delivered product to show that it is acceptable for use. Certification is a broader authoritative determination that depends on the domain, organization and applicable rules. In this course, those terms must not be used as synonyms.

For Mars hardware, a qualification campaign might expose a representative design to temperature, vibration, dust or other stresses at defined levels. Acceptance on each delivered unit might then verify workmanship and performance without repeating every qualification extreme. The system-level mission review still needs configuration traceability: evidence from one hardware revision cannot silently approve a later revision whose critical design changed.

8. End-to-end thread tests reveal failures that component tests cannot

A thread test follows one operational function across every subsystem required to complete it. "Receive medical telemetry" may require a sensor, local network, crew procedure, storage service, communications terminal, relay, ground link, software decode and human interpretation. Testing each box separately cannot prove that the chain works as a system.

1. **Communications-loss thread.** Remove the Earth link, verify local time/state knowledge, crew authority, onboard procedures, delayed-event handling and safe reconnection.

2. **Cabin-leak thread.** Inject a leak indication, test detection, localization, crew alarm, valve/ isolation authority, oxygen accounting, refuge transition and post-event configuration control.
3. **Power-deficit thread.** Remove a major generation source, exercise automatic load shedding, manual priorities, battery endurance, thermal consequences and recovery sequencing.
4. **Rover rescue thread.** Strand a vehicle near the protected radius and follow position knowledge, communications, crew consumables, rescue dispatch, charging/towing interfaces and return margin.
5. **Critical-spare thread.** Fail a declared line-replaceable unit and demonstrate diagnosis, physical access, tool availability, part identity, software/configuration update and functional return-to-service.
6. **Handover thread.** Give a fresh crew the released mission package and require them to reproduce configuration, protected margins, open waivers, known degraded states and the next decision gates without relying on oral memory.

9. Change control prevents a local fix from becoming a system regression

Every accepted change should identify what requirement, interface, hazard, procedure, model, verification result and operations product it can affect. This is especially important in a settlement architecture where software updates, locally manufactured parts and repair substitutions may occur far from Earth. “It works” is not enough; the mission needs to know what configuration now exists and which evidence remains applicable.

A practical change record therefore includes the reason, exact configuration delta, affected interfaces, updated analyses, required retest, temporary limitations, rollback path and approving authority. Emergency changes may use an expedited path, but they still need retrospective configuration capture. Otherwise the next failure begins with the team not knowing what system it actually owns.

10. Risk registers must expose dependencies and evidence, not multiply convenient probabilities

A risk statement should connect cause, event and consequence, identify the affected mission objective, and name mitigation with an owner and evidence. Common-cause dependencies prevent simplistic multiplication of subsystem reliabilities. Two redundant pumps sharing one contaminated feed, one software image or one cooling loop are not independent protection against those shared causes.

The final mission board should therefore keep dependency maps beside risk rankings. A low numerical probability with weak evidence can deserve more attention than a higher probability with excellent detection and recovery. Risk retirement requires evidence that the cause or consequence is controlled; moving a cell from red to green because the team has discussed it is not retirement.

11. R75 integrated 500-sol verification campaign — teaching sequence

Campaign gate	Injected condition	Evidence required before proceeding	Protected decision
Gate A — pre-departure	One critical subsystem carries only minimum accepted margin.	Trend review, open-waiver rationale, verified spare/ recovery path, no hidden reserve borrowing.	Launch, delay or de-scope.

Campaign gate	Injected condition	Evidence required before proceeding	Protected decision
Gate B — cruise	Communications outage overlaps a maintenance action.	Autonomy horizon, command authority, onboard procedures, protected power/ consumables.	Continue maintenance, defer it or enter safe configuration.
Gate C — first 30 sols	Power generation is 18% below planning value for several sols.	Updated energy ledger, thermal consequences, battery cycling impact, load-shed priorities.	Reduce science/industry loads or consume reserve.
Gate D — mid-campaign	One water-recovery component fails and its replacement consumes the only spare of that type.	New water horizon, repair evidence, next-failure exposure, resupply/ manufacture options.	Restore nominal operation or shift to conservation mode.
Gate E — rescue event	A rover is immobilized near the protected radius during poorer-than-planned communications.	Position confidence, crew consumables, rescue vehicle energy, route, weather/environment assumptions.	Dispatch, shelter in place or reduce rescue exposure by another tactic.
Gate F — Sol 500 handover	Incoming crew must assume control with two accepted degraded states.	Configuration baseline, open risks, waivers, margins, spare inventory, anomaly chronology, next verification actions.	Accept handover or require closure before authority transfer.

12. Final board dossier — twelve artifacts that must agree with each other

The capstone is closed only when the following products describe the same architecture and configuration: mission objectives and stakeholder expectations; requirements tree; interface register; mass/power/thermal/data/consumable budgets; risk and hazard register; technical performance measures and margin trends; verification matrix; validation plan; configuration/change log; qualification and acceptance evidence; operations concept with decision authority; logistics/spares ledger; and the final handover package. Contradictions between these artifacts are findings, even when every document looks polished in isolation.

The learner's oral defense should therefore be adversarial. Reviewers can select any protected requirement and ask for its parent need, implementation, interface dependencies, verification method, evidence, current configuration and remaining margin. If the answer requires searching several disconnected spreadsheets whose assumptions disagree, the mission is not closed.

Primary systems-engineering bridges

- NASA Systems Engineering Handbook
- NASA SE Handbook appendices — requirements, verification and validation matrices
- NASA-HDBK-1009A — Systems Modeling Handbook for Systems Engineering

R75 closure rule for module 16. The mission is not “integrated” until requirements, interfaces, margins, risks, verification, validation, configuration and operations all point to the same released architecture.