

Capstone project: design a 30-person Mars settlement



Conceptual visualization — the 30-person capstone must be solved as one coupled system rather than as independent lists of habitat, power, water and logistics equipment.

Integrate propulsion, power, habitat, water, food, mobility, maintenance, medicine and operations into one quantified and defensible architecture. This capstone project deliberately combines multiple disciplines in one architecture so tradeoffs cannot be hidden.

Before starting — review integrated mission architecture, systems engineering and logistics because this capstone accepts a settlement only as a coupled system. Every important symbol is defined at first use.

Mastery objectives

- integrate habitat, utilities, logistics, health, mobility and maintenance into one 30-person settlement architecture
- distinguish installed, available, verified and N-1 capacity before admitting crew
- build commissioning gates that prove interfaces and degraded modes rather than only individual equipment
- defend a settlement-wide GO/HOLD/NO-GO decision with evidence, human workload and recovery paths

1. A capstone is not a summary

This project forces decisions to agree with each other. An excellent power solution may be incompatible with logistics mass; an ambitious crop area may create unfunded electrical and thermal loads; a lightweight architecture may lack redundancy. The goal is interface management and trade studies. Every important assumption is stated, every major number is traceable and every margin has a reason.

Moon to Mars Architecture provides a framework for connecting habitat, logistics, mobility, power and other functions in one architecture. [NASA — Moon to Mars Architecture](#)

2. Define mission, crew and time horizon


Start with scenario: 30 people, nominal stay, resupply windows, return capability, growth phase and target autonomy. Without that definition, stocks, habitat and industry cannot be sized.

Separate what must operate before crew arrival from what can be assembled by crew and what

belongs to later expansion.

3. Mass budget and predeployment

Hardware does not arrive in one shipment. Build a manifest for habitat, power, ECLSS, consumables, rovers, workshop, spares and reserves. Critical elements are predeployed and tested when possible. A plan that requires one cargo vehicle to arrive successfully just days before crew has excessive logistics risk.



A 30-person settlement is accepted as a coupled system; one unverified interface can invalidate several apparently healthy subsystems. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

4. Energy and power budget

Separate daily energy from instantaneous power. ECLSS, computing and safety are continuous loads; workshop, ISRU and some food operations can be scheduled. Peaks must fit generation and storage. Define a degraded mode in which only vital loads remain powered.

5. Water, atmosphere and food

Life-support loops use realistic recovery and loss rates. Water combines recovery, storage and make-up source. Oxygen combines stores and production. Food combines imported reserves and local crops. Every loop has a survival time if production stops so that a headline “98% recycling” cannot hide an undersized reserve.

6. Habitat, fire and refuge

Show compartments, airlocks, refuge, evacuation routes, noisy zones, laboratory, workshop and storage. Fire and pressure barriers must not sever both routes to refuge. Cable, fluid and ventilation interfaces are identified because they can propagate failures across compartments.

7. Mobility, science and outside work

The settlement has a fleet with rescue capability. Routes to ISRU, science and landing sites are documented. EVA radius respects rescue and environmental constraints. Science operations preserve context and chain of custody without blocking daily maintenance.

8. Maintenance, spares and local industry

Create a critical-equipment list, redundancy strategy and spares policy. The workshop can make simple parts and refurbish components but does not instantly reproduce Earth's whole industrial base. A make/repair/stock/import matrix exposes true dependencies.

Deepening: success and abort criteria

A complete architecture defines when the mission can continue, when objectives must be reduced and when return or evacuation becomes necessary. Criteria are tied to measured resources such as days of water, power capability, ECLSS state, medical status, fleet and communications. Writing them before a crisis reduces improvisation and reveals where additional redundancy is worth its mass.

Deepening: coherent margins across subsystems

Adding 20% margin everywhere can make an architecture unnecessarily heavy, while inconsistent margins can hide the true weak point. Margins follow uncertainty and criticality: immature hardware mass, vital reserves, variable loads or ISRU performance. A system review checks that assumptions used by one subsystem match assumptions used by the others.

Deepening: Earth-dependency matrix

The final design distinguishes locally autonomous functions from those still dependent on Earth. For every critical function, list parts, software, consumables, skills and materials that cannot yet be replaced locally. This prevents calling a settlement autonomous because it produces oxygen while one catalyst, sensor or drug has no replacement path. The matrix ranks dependencies and identifies which reduction gives the largest resilience gain.

9. Worked example: water reserve under degraded mode

Thirty people at 3.5 L/day potable and food water require 105 L/day. At 97% recovery, theoretical loss is 3.15 L/day, so 1,000 L could cover over 300 days of that nominal loss. If recovery falls to zero, the same stock lasts only $1,000 \div 105 \approx 9.5$ days. Emergency sizing must use the outage case, not only nominal efficiency.

Calculated case study: makeup water for a 30-person settlement

TEACHING ASSUMPTION — Thirty people each use 12 L/day inside the accounting boundary. The loop recovers 92% of the water. Size 60 days of makeup water with a 20% reserve.

Let N be number of people; q gross requirement in L/person/day; η recovery fraction, dimensionless; d duration in days; and V volume in litres.

Daily gross requirement: $V_b = N \times q = 30 \times 12 = 360 \text{ L/day}$. Makeup fraction = $1 - 0.92 = 0.08$. Daily makeup = $360 \times 0.08 = 28.8 \text{ L/day}$. Over 60 days: 1,728 L. With a 20% reserve: $1,728 \times 1.20 = 2,073.6 \text{ L} \approx 2.07 \text{ m}^3$.

A 92% loop greatly reduces makeup stock but does not eliminate it. The 12 L/day value is a teaching assumption limited to this accounting boundary.

10. Integrated exercise

Design a 30-person architecture with 500 kW average power, 3 MWh storage, two pressurized habitats, crop area, workshop and three rovers. Define priority loads, 30-day reserves, redundant functions and actions if 40% of power generation is lost for ten sols.

11. Reasoned solution

Protect ECLSS, thermal control, communications, medical capability and safety first. Industrial loads and some artificially lit agriculture can be reduced or shifted. Storage should bridge peaks rather than be exhausted in the first sol. State the criteria for returning to normal and the effect of degraded operation on reserves.

12. Final design package

Deliver functional architecture, mass budget, energy/power budget, water, atmosphere, food, thermal system, maintenance, fleet, risks, emergency procedures, sources and assumptions. A critical review should identify the five largest remaining Earth dependencies and the three changes offering the best resilience gain per kilogram imported.

Capstone: design a thirty-person Mars settlement as one integrated system

This capstone is not a summary of previous lessons. It is an integration exercise in which every subsystem constrains the others. A thirty-person settlement needs habitat volume, power, thermal rejection, water, atmosphere control, food, mobility, maintenance, medical capability,

communications, governance, spares and an expansion plan. A design is credible only when its assumptions are explicit and its interfaces close.

The central discipline is consistency. A power system sized for normal operation but not for water processing after a failure is incomplete. A food plan that depends on greenhouse lighting must appear in the electrical and thermal budgets. A rescue rover is not available if the maintenance plan assumes it is routinely used for cargo. The capstone therefore uses one master assumptions table and one dependency matrix so changes propagate across the design.

The exercise also separates architectural choices from facts. NASA and other agencies provide mission architectures, research and standards, but this thirty-person settlement is a Delta-Sierra systems exercise. Numerical assumptions are planning values to be tested, not claims that an agency has selected this exact architecture.

1. Define mission, population and time horizon

Write down who the thirty people are, how long they must operate before the next resupply opportunity, what level of scientific and industrial activity is expected, and whether the settlement is a temporary outpost or the first phase of long-duration growth. These choices affect every stock and capacity calculation.

Population is not only a headcount. Skills, shift coverage and medical redundancy matter. If only one person can maintain a critical power converter, the system has a human single point of failure even when the hardware is redundant. The crew matrix should therefore identify minimum competencies per shift and cross-training requirements.

$$\text{Need_total} = N \times \text{need_per_person}$$

This simple scaling relation is useful for water, food, sleeping volume and many consumables, but it should not be applied blindly to shared infrastructure. Some systems scale with population, some with peak concurrent use, and some have fixed overhead.

2. Phase cargo before crew arrival

A robust architecture prepositions the systems that people need immediately: power generation and storage, communications, verified landing or unloading zones, initial habitat pressure integrity, water and atmosphere capability, emergency refuge, tools and a minimum mobile asset. Robotic commissioning can reveal failures before lives depend on the equipment.

Every cargo has an activation dependency. A power unit may need deployment equipment; a habitat may need a pressure source; a rover may need charging; a greenhouse may need water treatment. The cargo plan should therefore be a directed network rather than a list of masses. No critical node should depend on equipment arriving later than the node itself.

3. Habitat, zoning and refuge

The settlement should separate sleeping, hygiene, food preparation, medical care, workshop activities and potentially contaminating industrial work. Zoning reduces cross-contamination and helps isolate failures. Airlocks and suit interfaces belong at the edges of occupied zones so dust does not migrate unnecessarily into clean areas.

A refuge is a pressure-safe, fire-separated place where the crew can survive a serious habitat incident. It needs independent or isolatable atmosphere control, communications, emergency power and enough water and consumables for the planned response time. A refuge that depends on the same failed loop as the main habitat is not independent.

4. Water, atmosphere and food budgets

Water accounting should distinguish potable, hygiene, food, greenhouse, process and contingency uses. Recovery systems reduce imported demand but do not eliminate losses. Storage should support both normal smoothing and degraded operation when a processor is offline.

Atmosphere control tracks oxygen generation or storage, carbon-dioxide removal, humidity and trace contaminants. The critical question is not only average capacity but how long the crew can remain safe after loss of one major processing train. Repair time therefore appears directly

in storage requirements.

Food combines imported reserve and demonstrated local production. The capstone should show protected emergency stock and identify which nutrients remain Earth-dependent. Local agriculture is credited only at a production level demonstrated by the chosen equipment and operating assumptions.

5. Microgrid: power generation, storage and priorities

The electrical architecture should identify critical, essential and discretionary loads. Critical loads include functions whose interruption threatens life or immediate safety. Essential loads support continued mission capability but may tolerate controlled interruption. Discretionary loads such as some manufacturing tasks can be shed when generation is constrained.

$$\text{Capacity_required} = \text{Demand_peak} \times (1 + \text{margin}) / \text{availability_factor}$$

If peak critical-and-essential demand is 120 kW, planning margin is 25% and the chosen availability factor is 0.8, required installed or dispatchable capacity in this simplified planning relation is $120 \times 1.25 / 0.8 = 187.5$ kW. The equation is not a substitute for time-series simulation, but it reveals the cost of relying on hardware that is not always available.

Energy storage must be sized for duration as well as power. A battery that can deliver 200 kW for fifteen minutes is different from a system that can support 80 kW for ten hours. The capstone should include at least one credible low-generation or fault scenario and show how loads are shed over time.

6. Thermal architecture follows the same load hierarchy

Power consumed inside the settlement becomes heat that must be transported or rejected. The thermal design should show normal and degraded heat-rejection capacity, cross-ties between loops, heat recovery and temporary storage. High-power industry may be scheduled when both electrical and thermal margin are available.

Habitat survival should not depend on an industrial cooling loop that cannot be isolated. Shared infrastructure can reduce mass, but interfaces must permit emergency separation. The capstone should state which loops can be isolated and what capacity remains after isolation.

7. Mobility and rescue geometry

Define the routine logistics network, crew exploration radius, landing-zone connection and emergency rescue envelope. A thirty-person site may need multiple vehicle classes rather than one universal rover. Cargo vehicles can prioritize payload and autonomy while crew vehicles prioritize life support, redundancy and recoverability.

The rescue plan must account for simultaneous tasks. If two crews operate in different directions, the settlement needs enough ready mobility to respond to either without abandoning the other. Rescue readiness is therefore a scheduling constraint, not merely a vehicle inventory line.

8. Maintenance, spares and first local industry

Maintenance converts hardware reliability into settlement availability. The plan should identify line-replaceable units, diagnostic tools, calibration needs, lubricants, seals, filters, connectors and the inventory of high-consequence spares. Parts with long Earth lead times need stronger local stock than common low-consequence hardware.

Local manufacturing begins with items that are feasible to make and valuable to replace: simple brackets, ducts, covers, fixtures, cable management, some structural parts and repair aids. It should not be credited as the ability to reproduce every complex electronic, medical or pressure-critical component. The dependency matrix keeps those limitations visible.

9. Medical capability and continuity of care

The medical plan starts from plausible clinical scenarios and the time to Earth resupply or evacuation, not from an arbitrary medicine count. Diagnostics, sterile consumables, oxygen, temperature-controlled stock and trained personnel form one system. Some equipment is useless if the required consumables or competence are missing.

Cross-training reduces human single points of failure. The capstone should show how many crew can perform first response, advanced care, equipment maintenance and medical decision support. Telemedicine helps but cannot remove communication delay or substitute for local stabilization capability.

10. Science and industrial work compete for shared resources

Science laboratories require clean power, controlled environments, sample handling and crew time. Industrial processes can create dust, vibration, chemical hazards or high thermal loads. Zoning and scheduling should prevent routine industrial work from compromising scientific measurements or habitat safety.

The settlement should define its first industrial priorities based on dependency reduction. Producing a locally useful reagent, construction material or spare can be more valuable than maximizing total tonnes processed if it removes a critical Earth dependency.

11. Governance and human workload

A small isolated settlement needs clear operational authority without making every decision rigid. Roles should distinguish routine operations, technical authority, medical authority and emergency command. Decision logs are valuable because they preserve why a configuration changed and help later crews understand past tradeoffs.

$$\text{Workload_ratio} = \Sigma \text{ task_hours} / \text{available_crew_hours}$$

If required scheduled work totals 540 person-hours per week and available planned work capacity is 720 person-hours after sleep, meals, exercise, personal time and reserve, workload ratio is 0.75. A design that routinely requires a ratio near 1.0 has little capacity for repairs, emergencies or learning and is therefore fragile even if every engineering subsystem is nominally adequate.

12. Earth dependence and commissioning

Earth dependence should be measured by category. Food may be partly local while electronics remain almost entirely imported. Medical consumables, catalysts, specialty polymers and software support can each have different dependency profiles. A single “percentage self-sufficient” number can hide the component that actually limits survival or growth.

$$\text{Earth_dependency} = \frac{\text{Imported_critical_mass}}{\text{Total_critical_consumed_mass}}$$

This ratio can be calculated for selected categories, but mass alone is not enough. A ten-gram sensor can be more critical than a tonne of bulk material. The capstone therefore pairs mass dependence with criticality and lead time.

Commissioning is the process of proving that installed systems operate as intended before full reliance. The schedule should test power, pressure integrity, atmosphere, water, communications, refuge and emergency modes before population rises to the target thirty people.

Master budget table for the capstone

Domain	Primary quantity	Normal requirement	Degraded-mode proof	Evidence to retain
Population	people and competencies	30 residents with shift coverage	minimum staffing after illness or injury	competency matrix
Power	kW and kWh	peak plus margin	critical loads after one major source or branch loss	load schedule and one-line diagram
Thermal	kW heat rejection	peak heat after recovery	capacity after loop or radiator derating	heat-flow map
Water	kg/day and stored kg	net use after recovery	days of operation with processor offline	water balance
Atmosphere	kg/day gases and removal capacity	steady crew metabolism	storage or backup after processor loss	atmosphere balance
Food	kcal/day plus nutrients	normal diet	greenhouse outage reserve	stock and crop table
Mobility	ready vehicles and kWh	daily logistics plus sorties	credible rescue while another rover is unavailable	dispatch map
Maintenance	spares, tools, person-hours	scheduled and corrective work	repair of high-consequence failures	spares criticality register
Medical	clinical functions and days coverage	routine care	stabilization during serious case	medical capability matrix

Reference solution logic: six injected events

Event 1 - One power source is lost during a high-load work period

Automatic or procedural load shedding protects life support, communications, water circulation and medical refrigeration. Industrial work stops first. Operators compare remaining generation and storage against the critical-load curve, then decide whether to reconfigure the microgrid or repair the failed source. The event is passed only if the settlement can show hours of protected operation, not merely that “backup power exists.”

Event 2 - Water processor is unavailable for forty-eight hours

The crew switches to stored water and reduces non-essential use. Tank inventory, expected repair time and minimum hygiene requirements become the decision variables. If storage was sized only for normal hourly smoothing, the architecture fails this event even if average water recovery efficiency is excellent.

Event 3 - A crew rover becomes immobile beyond walking distance

The rescue schedule confirms that another crew-capable vehicle is ready and has enough energy, payload capacity and life-support margin. Cargo movements are delayed if necessary. The settlement passes when rescue capability survives routine logistics rather than being consumed by them.

Event 4 - Greenhouse output falls sharply for two months

Food operations quantify the deficit by calories and limiting nutrients, release protected stock according to policy and prioritize greenhouse recovery. The power budget may also change because damaged growing capacity uses less lighting while repair equipment creates new loads.

Event 5 - A critical imported spare fails with no replacement on Mars

The maintenance team examines cannibalization, derated operation, local fabrication of non-critical subparts and configuration changes. If the component cannot be reproduced, the design should already have redundancy or a stocked replacement. The event exposes false claims of self-sufficiency.

Event 6 - Several crew members are unavailable at the same time

The crew matrix is recomputed. Some science or construction activity stops so that power, life support, medical care and maintenance remain staffed. A technically perfect settlement that requires every specialist every day has no human resilience.

Capstone calculations

Water reserve for processor downtime

$$\text{Reserve_water} = N \times \text{net_use_per_person} \times \text{outage_days} \times (1 + \text{margin})$$

For an illustrative net contingency use of 12 kg/person/day, thirty people, a three-day outage and 25% margin, reserve is $30 \times 12 \times 3 \times 1.25 = 1,350$ kg. This is a planning example; the real value must be tied to the chosen water architecture and operational conservation rules.

Food reserve

$$\text{Energy_reserve} = N \times \text{kcal_plan} \times \text{days} \times (1 + \text{margin})$$

At an illustrative 3,000 kcal/person/day for 30 people and 45 days, baseline is 4,050,000 kcal. With 15% planning margin the reserve target is 4,657,500 kcal. Nutrient balance and food stability must be checked separately.

Mobility readiness

$$\text{A_mobility} = \frac{N_mission_ready}{N_required_for_safe_operations}$$

If safe simultaneous operations require four mission-ready vehicles and only three are available, the ratio is 0.75 and the schedule must be reduced. Counting parked but unserviceable vehicles would hide the real constraint.

Commissioning duration

$$t_{\text{commissioning}} = \max(t_{\text{critical_paths}})$$

The settlement cannot be declared ready before the longest safety-critical dependency chain is complete. If power commissioning takes 5 days, water 8 days and pressure/atmosphere verification 11 days, the minimum commissioning time is at least 11 days unless tasks are reorganized without compromising required sequence.

Deliverables required from the learner

The final dossier should include a mission statement, assumptions register, cargo phasing, settlement plan, power and thermal budgets, water/atmosphere/food balances, mobility and rescue map, maintenance and spares register, medical capability matrix, staffing and governance plan, Earth-dependency matrix, commissioning sequence and responses to the six injected events. Every major numerical claim should state its source or planning status.

A strong submission makes uncertainty visible. It distinguishes measured values, sourced design references and local planning assumptions. It shows what happens after failures and how the crew knows when to stop an activity. The result should read like an operations-ready concept, not a promotional vision.

Progressive exercises with solutions

- ▶ Exercise 1 - Power capacity
- ▶ Exercise 2 - Water contingency
- ▶ Exercise 3 - Workload
- ▶ Exercise 4 - Critical-path commissioning

Interactive beginner glossary

- capstone - integrated final systems project.
- assumption - planning input that is not automatically a fact.
- critical path - dependency chain that sets schedule duration.
- commissioning - verification before service.

- degraded mode - safe operation with reduced capability.
- single point of failure - one failure that can remove a critical function.
- dependency matrix - structured record of interfaces and dependencies.
- resilience - capacity to continue through failures.

Capstone method: design the thirty-person settlement from functions before hardware

A capstone becomes useful when it forces the student to integrate the whole curriculum. The design should therefore begin with required functions rather than favorite technologies. A settlement needs breathable atmosphere, safe water, food, power, thermal control, medical capability, communications, mobility, maintenance, waste handling, shelter, fire protection, governance and a way to recover from failures. Hardware is selected only after each function has a measurable requirement and a failure consequence.

Step 1 — define mission states before sizing equipment

At minimum, define nominal, degraded, emergency and recovery states. A nominal state describes ordinary operations. A degraded state preserves essential functions with reduced margin. Emergency state protects life and habitat integrity. Recovery state restores barriers, reserves and documentation after stabilization. Sizing only for nominal conditions produces a settlement that looks efficient on a spreadsheet but has no coherent behavior when something breaks.

For each state, list which functions must continue, which may be reduced, and which may stop. A greenhouse light schedule might be reduced during a short power emergency; atmosphere circulation, fire detection or medical support may not have the same flexibility. The design should explain these distinctions rather than hiding them inside one average power number.

Step 2 — write a requirement as a measurable sentence

“Provide enough water” is not a requirement. A useful requirement names quantity, quality, duration, condition and verification method. The same applies to power, communications and medical capability. The student should be able to point to the sensor, record or test that proves each requirement is being met.

Requirements should also expose uncertainty. If local water extraction is expected but not guaranteed at a particular rate, the settlement must either carry enough reserve to survive a lower production case or have an alternative source. The uncertainty belongs in the architecture, not in a footnote added after sizing.

Step 3 — build a dependency matrix

Every major function should be placed both as a row and a column. Mark when one function depends on another. Water processing may depend on electrical power, pumps, sensors, chemistry and laboratory verification. Food production may depend on water, power, nutrients, thermal control and human labor. Medical capability may depend on power, cold storage, communications, clean water and trained personnel. A dense matrix reveals why losing one apparently small shared utility can disable several independent-looking systems.

Common-cause analysis then asks whether nominal redundancy is truly independent. Two water processors on the same bus, two oxygen sensors using the same software, or two rovers requiring the same unavailable spare are not equivalent to two independent protections.

Simple dependency exposure indicator

$$D_{\text{common}} = \frac{N_{\text{functions_using_common_resource}}}{N_{\text{critical_functions}}}$$

This teaching ratio does not calculate risk. It simply forces the designer to count how many critical functions share one resource. If six of ten critical functions depend on one electrical bus, $D_{\text{common}} = 6/10 = 0.60$. The correct question is then whether that common bus has adequate protection, alternative routing and repair strategy.

Step 4 — design buffers as time

Mass and energy stores should be converted into time to consequence. Ten tonnes of water sounds large, but its meaning depends on net daily loss in the current operating state. A battery bank sounds impressive until the protected load is applied. Spare filters matter only when their

consumption or failure rate is connected to a replacement timeline. Expressing buffers in hours or days makes cross-system decisions easier during a crisis.

Step 5 — allocate maintenance, not only equipment

A thirty-person settlement has a finite pool of attention. Every pump, valve, suit, rover, sensor and software service creates inspection, calibration, cleaning, repair and documentation work. The capstone should therefore contain a maintenance workload budget. If the proposed architecture requires more technician-hours than the crew can supply while also performing science, health, food production and operations, the design is not closed.

Maintenance workload fraction

$$F_{\text{maint}} = H_{\text{maintenance}} / H_{\text{technical_available}}$$

If the technical crew can provide 420 person-hours in a planning period and scheduled plus expected corrective maintenance requires 315 person-hours, $F_{\text{maint}} = 315/420 = 0.75$. That leaves 25% for unexpected work, upgrades and technical support. A ratio near or above one means the plan consumes all available technical attention before surprises occur.

Step 6 — commission the settlement in proof gates

Do not imagine that thirty people arrive and every system starts at once. A robust deployment sequence proves power, communications, environmental control, water, fire protection, refuge, logistics and maintenance capability before population and industrial demand rise. Each gate should have objective entry and exit criteria. A failed test pauses progression and creates a corrective action rather than being waived by schedule pressure.

Integrated capstone drill

Assume one power converter fails, reducing available generation; a water-quality sensor simultaneously drifts; and one pressure suit is removed from service. The student must identify which functions are immediately threatened, what information is trustworthy, which loads are

shed, how many EVA tasks are deferred, how water quality is independently checked, what protected reserves are consumed, and what evidence permits return to nominal state. The solution should include a timeline, not just a list of actions.

The capstone is successful only if the proposed settlement can explain how it behaves when assumptions are wrong. A beautiful nominal architecture is not enough. The student should be able to trace every life-critical function from requirement to sensor, actuator, backup, spare, trained person, procedure and recovery criterion.

Thirty-person capstone laboratory: every subsystem must expose its survival margin

The capstone is not complete when the equipment list is complete. Each critical function needs a measurable reserve, a failure clock and a recovery route.

Water buffer time

$$D_{\text{water}} = m_{\text{stored,usable}} / m_{\text{net-loss,daily}}$$

Teaching scenario. If 1,200 kg of usable contingency water is protected and the degraded settlement loses 60 kg/day net after all functioning recovery processes, buffer time is $1,200/60 = 20$ days. If net loss doubles, buffer time halves.

Interpretation. This is a response clock. It does not say the crew should wait 20 days to act; it tells planners how much time exists before the protected stock is exhausted under the stated state.

First-pass power scaling

$$P_{\text{base}} = P_{\text{fixed}} + N \times p_{\text{variable}}$$

Suppose a teaching model assigns 50 kW fixed settlement infrastructure plus 1.5 kW average variable demand per resident. For 30 residents: $50 + 30 \times 1.5 = 95$ kW. Adding 25% planning margin gives $95 \times 1.25 = 118.75$ kW. This is deliberately simple: real loads are time-varying and some do not scale linearly with population.

► Exercise — simultaneous margin loss

DELIVERY

mass · spares ·
feedstock

HABITAT

pressure · refuge ·
zoning

UTILITIES

power · water ·
thermal

ACCEPTANCE

measured margins ·
go/no-go

Teaching composite: commissioning a 30-person settlement means proving a dependency chain, not merely delivering hardware.

Commissioning studio: prove a thirty-person settlement before thirty people depend on it

A settlement architecture is not complete when drawings look coherent. It becomes credible when every life-critical function can be commissioned, challenged and shown to recover from realistic faults before the crew population relies on it. Commissioning means moving from “installed” to “demonstrated under measured conditions.” The distinction is especially important on Mars because a subsystem that merely powers on may still fail under sustained load, dust exposure, simultaneous demand or loss of an upstream dependency.

Write acceptance evidence before launch

For each critical function, define an acceptance statement that contains a measurable quantity, an operating duration, a configuration and a failure test. “Water system works” is not an acceptance criterion. “The water loop supplies the protected daily demand for 72 hours while one processing train is unavailable, maintains quality within the selected limits and restores nominal configuration without violating the reserve floor” is much closer to one. The same logic applies to air revitalisation, power, thermal rejection, communications, food storage, fire refuge, medical capability and surface rescue.

Commissioning evidence margin

$$M_{\text{test}} = (C_{\text{tested}} - D_{\text{required}}) / D_{\text{required}}$$

Question. How much demonstrated capability exists above the requirement?

Symbols. C_{tested} is capacity actually demonstrated in the relevant test; D_{required} is the required capacity under the defined mission state; the ratio is dimensionless.

Example. A protected water train demonstrates 420 L/day while the degraded-state requirement is 360 L/day. Margin is $(420-360)/360 = 0.167$, about 16.7%.

Interpretation. The number is meaningful only if the test configuration matches the mission configuration. A 17% margin measured with clean filters and unlimited electrical power may disappear after months of operation.

Limit. One margin does not prove reliability. Duration, sensor validity, maintenance history and common-cause vulnerabilities still matter.

Crew time is a resource budget

Thirty residents do not provide thirty full-time engineers. Sleep, exercise, meals, hygiene, medical care, science, administration and training consume time. Maintenance demand can therefore become a hidden design constraint. A system that saves launch mass by requiring constant manual attention may be a poor settlement system.

Maintenance workload fraction – growth-gate application

$$F_{\text{maint}} = H_{\text{maint}} / H_{\text{available}}$$

If the crew can safely allocate 420 person-hours per week to technical operations and scheduled plus corrective maintenance consumes 190 hours, $F_{\text{maint}} = 190/420 \approx 0.45$. Almost half of the technical labour envelope is already consumed. A new process requiring another 90 hours per week would raise the fraction to $280/420 \approx 0.67$ before any emergency occurs.

Decision use. The calculation forces the designer to ask whether automation, better access, modular replacement or an additional specialist should be added before population growth.

Spare parts should be linked to consequences

A capstone inventory should not simply list thousands of components. Rank spares by time to consequence, probability of demand, commonality, repairability and whether the part can be fabricated locally. A low-mass seal that can stop an oxygen compressor for months may deserve more protection than a heavy non-critical panel. Conversely, carrying five complete replacements for a repairable component may waste mass that would have purchased broader resilience elsewhere.

The learner should build a critical-spares register that identifies the failed function, detection method, isolation action, interim configuration, repair skill, tools, consumables, expected repair duration and evidence required before return to service. That register connects logistics to reliability rather than treating stores as a separate warehouse problem.

Integrated commissioning sequence

Commissioning should proceed from dependencies upward. Verify electrical distribution and protected control power before relying on pumps; verify cooling before sustained high-power process tests; verify atmosphere monitoring before hazardous chemical work; verify communications and localization before long rover sorties. Then inject faults deliberately.

Isolate a pump. Remove one power source. Block a sensor input. Simulate a missed resupply. Demonstrate that operators recognize the state, enter the intended degraded mode and preserve protected reserves.

Capstone review — reject or accept crew arrival?

The habitat has passed pressure testing, but the water processor has only 18 hours of continuous demonstrated operation, one backup CO₂-removal unit has not been tested under load, and the rover rescue route has never been driven at night. Cargo windows make delay expensive. Write an acceptance decision with three categories: evidence sufficient for arrival, evidence that must be completed before arrival, and evidence that can be completed after arrival. Justify every category using time to consequence rather than schedule pressure.

► Reasoned solution

First-Man capstone: a 30-person settlement must earn permission to receive people

The capstone is not a drawing exercise. It is a readiness argument. A settlement for thirty people should be treated like a complex vehicle that will be occupied continuously and cannot rely on immediate outside rescue. Before crew arrival, the design team must convert every attractive subsystem into evidence that it works as part of the whole: power feeds life support, life support consumes spares, thermal control depends on power and fluid loops, logistics depends on mobility, medical capability depends on people and stocks, and every repair consumes crew time.

The most dangerous sentence in a design review is “that subsystem has enough capacity” when the word *enough* has not been tied to a boundary, duration and failure state. Installed capacity is not the same as commissioned capacity. A 200 kW power plant that has only demonstrated 120 kW continuously under the intended thermal and control configuration should not be credited as 200 kW in the readiness case.

Use a common margin language across unlike systems

$$M_j = (C_{j,\text{verified}} - D_{j,\text{design}}) / D_{j,\text{design}}$$

Starting question

How can the review compare margin in power, water processing, oxygen production or waste handling without pretending those capacities have the same units?

Read aloud

Read: "margin for function j equals verified capacity minus design demand, divided by design demand."

Symbols, pronunciation and meaning

$C_{j,\text{verified}}$ is the capacity actually demonstrated for function j; $D_{j,\text{design}}$ is the selected design demand; M_j is a dimensionless fractional margin.

Units

The numerator and denominator use the same unit for each function, so the ratio is dimensionless. Do not mix kW with kWh, kg/day with kg, or instantaneous flow with daily volume.

Origin and status of values

Verified capacity comes from commissioning evidence. Design demand comes from the agreed population, duty cycle and operating assumptions, not from a later optimistic reinterpretation.

Why this operation

Subtracting demand from verified capacity gives absolute spare capacity. Dividing by demand converts it to a comparable fraction.

Substitution and calculation

Suppose a water processor demonstrates 420 L/day while design demand is 350 L/day. Margin = $(420-350)/350 = 0.20 = 20\%$.

Calculator entry

Enter $(420-350)\div350\times100$ if you want percentage margin.

Mental estimate

Seventy extra litres over 350 is one fifth, so 20% is plausible.

Independent check

$350 \times 1.20 = 420$ L/day, recovering the verified capacity.

Physical or operational interpretation

A positive margin means the demonstrated capacity exceeds the chosen demand in that tested configuration. It says nothing yet about redundancy or repair time.

Plain-English translation

The processor has demonstrated about one fifth more daily capacity than the design case requires.

Variation / sensitivity

If real demand rises to 390 L/day, the same verified capacity provides only about 7.7% margin. Population and operating behaviour can quickly consume apparently comfortable headroom.

Limit / assumption

Margins across different functions cannot simply be averaged. A 50% power margin cannot compensate for a negative oxygen margin.

What this does not prove

The equation does not prove resilience to a failed pump, contaminated loop, blocked filter or unavailable technician. Those need separate degraded-mode evidence.

Boundary case to test

If C_{verified} is below demand, M is negative and the crew-arrival gate fails unless another independent capacity or a lower verified demand closes the deficit.

Build a commissioning matrix, not a completion percentage

A statement such as “the base is 93% complete” is almost meaningless. Readiness should be organised by critical function and evidence state: installed, powered, checked locally, tested under load, tested with interfaces, tested in degraded configuration, accepted, and supported by spares and procedures. A decorative interior can be unfinished while the settlement is safe; one unverified oxygen isolation valve can block occupancy.

Install

hardware present and identified

Commission

function demonstrated under real interfaces

Degrade

credible failure mode exercised safely

Accept

evidence, spares, crew and procedure ready

The hidden budget is crew time

A settlement can close every mass and energy balance and still fail because it requires more maintenance hours than the crew can supply. Estimate routine inspection, cleaning, filter changes, calibration, crop work, medical duty, software/configuration control, EVA preparation and unexpected repair. Then protect sleep, exercise, hygiene, training and scientific or productive work. A system that survives only by consuming all human time is not autonomous; it is converting crew life into maintenance margin.

Interfaces deserve their own acceptance tests

Failures often occur between subsystems. A water unit may meet its standalone flow rate but trip the power bus during a simultaneous high load. A rover charger may work but inject thermal load that the local loop cannot reject. A medical refrigerator may have emergency power but share a breaker with a nonessential load. The capstone review should therefore contain interface tests deliberately designed to make several systems interact at realistic peaks.

Thirty people require social as well as technical redundancy

For each critical competence, list who can perform the task independently, who can assist, and how long training a replacement would take. One expert in electrical protection, one surgeon, one greenhouse disease specialist or one software administrator can become a single point of failure even if the hardware is redundant. The pre-arrival readiness gate should expose these human singletons before they become emergency surprises.

Final readiness exercise

The base has positive nominal margins in power, oxygen, water and food. During commissioning, however, the backup water pump has never run under the final software configuration; the only technician able to service it is scheduled to arrive with the first crew; and the spare seal kit is still in an unlanded cargo vehicle. The correct decision is not “95% ready.” Water resilience is unverified and depends on a person and spare that are not yet present. Either commission the backup with available staff and parts, bring the missing capability forward, or delay occupancy.

Commissioning proves nominal and degraded behavior before population is admitted. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Thirty-person settlement commissioning dossier: prove capability before accepting the crew

A thirty-person Mars settlement is not ready because thirty bunks exist or because every subsystem has passed an isolated bench test. Readiness means that the integrated settlement can sustain the intended population through defined nominal and degraded states, with evidence for life support, energy, thermal control, food, habitat, mobility, medical response, maintenance, communications and human workload. The final capstone should therefore be reviewed as a **commissioning case**, not as a catalogue of installed equipment.

Separate installed, available and verified capacity

Installed capacity is what the hardware nameplates suggest. Available capacity is what can be operated in the current configuration. Verified capacity is what has actually been demonstrated under stated conditions. Population admission should be based on the third category for critical functions. A water unit rated for a certain throughput, for example, does not establish settlement capacity until pumps, storage, controls, sensors, distribution and recovery have operated together at the required load.

Build a master capacity table with one row per critical function and at least these columns: requirement, installed capacity, verified capacity, N-1 capacity where relevant, required reserve, current evidence, open anomaly, recovery path and owner. The table should make it impossible to hide a weak subsystem behind surplus elsewhere.

Verified population margin by critical function

$$M_{\text{pop},j} = (N_{\text{supported},j,\text{verified}} - N_{\text{planned}}) / N_{\text{planned}}$$

1 – Concrete question

For one critical function j , how much verified population capacity remains above the planned thirty residents?

2 – Intuition

Compare the number of people the demonstrated system can support with the population you intend to admit.

3 – Quantities

Determine supported population from verified capacity and the same demand assumptions used in the settlement design;

then compare with planned population.

4 – Formula

Population margin is supported population minus planned population, divided by planned population.

5 – Read aloud

“M pop j equals verified supported population for function j minus planned population, divided by planned population.”

6 – Symbols

j identifies a function such as oxygen, water, sanitation, food logistics or refuge capacity; N is a number of people.

7 – Pronunciation

The subscript j means the calculation is repeated independently for each critical function.

8 – Units

People minus people divided by people gives a dimensionless margin.

9 – Convention

Supported population must use the same reserve rule and operating state across functions; do not compare one nominal number with another contingency number.

10 – Why this relationship

The numerator is spare person-equivalent capacity; division by planned population expresses it as a relative margin.

11 – Assumptions

The conversion from hardware capacity to supported population is assumed valid for the stated demand and duty cycle.

12 – Unit check

$(\text{person-person})/\text{person} = 1.$

13 – Numerical case

Planned population: $N_{\text{planned}} = 30$ people.

Function A verified capacity: 42 people $\rightarrow$ margin = $(42 - 30)/30 = 12/30 = 0.40 = 40\%.$

Function B verified capacity: 35 people $\rightarrow$ margin = $5/30 = 0.1667 \approx 16.7\%.$

Function C verified capacity: 33 people $\rightarrow$ margin = $3/30 = 0.10 = 10\%$.

The limiting verified population margin is therefore 10%.

14 – Operations

Calculate each function separately. Refuge is the weakest in this example; averaging margins would conceal that fact.

15 – Algebra check

For refuge, $30 \times 1.067 \approx 32$ people.

16 – Mental estimate

Two extra refuge places above thirty is about one fifteenth, close to 6–7%.

17 – Interpretation

The settlement may have generous oxygen margin but only narrow refuge margin; the latter can control admission.

18 – What it does not prove

It does not prove all thirty people can reach the refuge in time, nor that refuge life support is independent of the initiating failure.

19 – Sensitivity

If one refuge compartment is unavailable and verified occupancy falls to 27, margin becomes negative: $(27-30)/30 = -10\%$.

20 – Practice

Guided exercise. Compute population margins for verified capacities of 42, 35 and 33 people against a planned population of 30.

Detailed guided correction.

1. 42-person capacity: $(42-30)/30 = 12/30 = 0.40 = 40\%$.
2. 35-person capacity: $5/30 = 0.1667 \approx 16.7\%$.
3. 33-person capacity: $3/30 = 0.10 = 10\%$.
4. The limiting margin is 10%, so the 33-person function is the first population bottleneck under the stated verified capacities.

Autonomous exercise. A 30-person settlement shows nominal verified capacities of oxygen 39, water 36, sanitation 34 and refuge 33. Under N-1 conditions they fall to 31, 32, 28 and 30. Build the N-1 margin column and decide whether crew arrival should be accepted.

► Autonomous correction – open after attempting the exercise

21 – Mission decision

Population admission is HOLD whenever any mandatory function lacks the required verified margin, even when the average system margin is positive.

Commissioning sequence: from dead hardware to inhabited settlement

The capstone should present an explicit sequence. First comes site power and communications sufficient for robotic work. Then storage, thermal rejection and core environmental systems. Next, pressurised volume is leak-tested, compartment isolation is demonstrated and fire/depressurisation response is rehearsed. Water and atmosphere loops are run under sustained load. Workshops and spares are commissioned before they are needed for an emergency. Mobility and rescue routes are tested. Medical receiving capability is stocked and exercised. Only after the settlement has accumulated stable operating evidence should the first permanent population increment be admitted.

Each stage needs entry criteria and exit criteria. “Water processor installed” is not an exit criterion. “Seven consecutive sols at defined throughput with acceptable quality, stable storage balance, alarm tests completed and one planned maintenance intervention recovered” is closer to the type of evidence a serious review needs. The exact criterion is mission-specific, but the form of evidence must be explicit.

Predeployment and dependency logic

Every crew-critical function should have the equipment, consumables, tools and spares needed for its expected early failures before the crew depends on it. This forces a cargo-phasing question: what must arrive one window earlier? A replacement pump shipped on the same flight

as the crew cannot protect commissioning that must occur before crew arrival.

Primary-source bridge. NASA DRA 5.0 provides historical human-Mars architecture context for predeployment, surface systems and operational concepts. [NASA NTRS — Human Exploration of Mars DRA 5.0](#).

Construct a dependency graph. Water processing may depend on power, thermal control, sensors, chemicals, filters, storage and maintenance tools. The workshop may depend on ventilation and electrical power. The medical system may depend on communications, refrigeration and clean water. If several nominally redundant systems depend on one shared utility, that shared utility can become a common-cause vulnerability.

Human workload is a capacity like power or water

Thirty residents create enough tasks to overwhelm a small number of specialists. The design should budget routine operations, maintenance, science, food production, EVA, exercise, medical duties, training and administrative work. Peak workload during faults matters more than average workload. A design that requires the same two people to troubleshoot power, restore thermal control and support an EVA rescue is not resilient even if every hardware subsystem has redundancy.

Primary-source bridge. NASA's Moon to Mars objectives explicitly identify crew time, maintainability and reuse as recurring architecture considerations. [NASA — Moon to Mars Strategy and Objectives](#).

For every critical discipline, list minimum qualified staffing, normal qualified staffing, backup personnel, recency requirements and the training path that turns a backup into a competent operator. Human single points of failure should appear in the same risk register as hardware single points of failure.

Six-event integrated review

The final dossier should survive at least six linked injections rather than six isolated textbook faults. Example sequence: a cargo pallet with filters is delayed; a dust event reduces solar generation; one thermal loop pump shows abnormal vibration; a crew member is temporarily

medically unavailable; a rover used for maintenance support is out of service; and a water-quality result becomes suspect. The exercise is not to “solve” all six independently. It is to protect the settlement priorities while resources, people and recovery paths interact.

The review team should demand a timeline. At minute zero, what is known? At hour two, which loads are shed? At hour twelve, which inventory is being consumed faster than planned? At sol two, what maintenance is deferred and what new risk does that create? A system that looks resilient in a static block diagram can reveal brittle dependencies once time is added.

Final evidence package

A serious thirty-person capstone should end with a compact but auditable evidence package: assumptions register; population and mission timeline; mass and cargo phasing; power and energy budgets; thermal balance; water/air/food inventories; habitat and refuge plan; mobility and rescue map; maintenance and spare strategy; medical capability; workforce/qualification matrix; dependency graph; top hazards; commissioning evidence; contingency modes; open risks; and explicit GO/HOLD/NO-GO criteria. Every headline claim should point to a calculation, test, source or clearly labeled assumption.

Qualification drill: refuse an attractive crew arrival

Your settlement can physically house thirty people and most systems are green. However, one refuge compartment has not completed leak verification, the workshop has only one qualified electrical maintainer, and a water-loop spare pump will arrive with the crew rather than before them. Write the review decision. Identify which items are merely inconvenient and which violate crew-admission logic. Then propose the smallest predeployment or training changes that would turn HOLD into GO without pretending the unresolved evidence is already closed.

Source context. NASA Moon-to-Mars architecture, ECLSS material and rover work provide system context. The integrated thirty-person design remains a Delta-Sierra teaching architecture, not an announced NASA mission. [NASA — Moon to Mars Architecture](#).

Reference architecture review: thirty residents as a network of functions

A useful reference solution can be described without pretending that one exact design is universally correct. Divide the settlement into at least two pressurised habitat/fire zones, protected life-support rooms, a workshop/industrial zone, food-production space, medical

receiving/isolation capability, storage and a set of external utilities. Every zone should have explicit pressure, power, data, ventilation and access dependencies. A fault in one zone must not automatically disable every refuge or every route to essential equipment.

At thirty residents, the architecture is large enough that shift structure matters. Some crew sleep while others operate machinery or conduct EVA. Alarms, maintenance isolations and emergency communications must therefore work across shifts. Quiet hours cannot imply reduced fire watch, and an off-duty specialist cannot be the only person able to reset a critical system.

Cargo phasing as a proof problem

Build a cargo table by function and dependency rather than by launch vehicle alone. For each critical item, identify when it must be on Mars relative to the first crew dependency.

Commissioning hardware, diagnostic tools and initial spares often need to arrive before the crew that will rely on them. A second identical component is not useful redundancy if both are delivered in the same high-risk cargo event or stored in the same vulnerable location.

Predeployment should also include consumables used during commissioning. Filters, calibration gases, cleaning agents, seals, sampling supplies and test loads may be consumed before nominal operations begin. The design should not arrive at “crew-ready” status with safety stock already spent proving the system.

Mass budget: protect uncertainty explicitly

Do not bury reserve mass inside subsystem estimates. Maintain a current estimate, growth allowance, uncertainty or maturity reserve and programme-level unallocated reserve. The categories serve different purposes. If a subsystem repeatedly consumes reserve because its definition is immature, that is a management signal rather than a reason to rename the margin.

Mass is coupled to volume, power, thermal control and landing architecture. A “small” addition can trigger a larger packaging or landing change. The capstone should therefore record second-order effects: extra batteries add mass and heat; more water adds mass but also shielding and thermal inertia; larger workshops need tools, ventilation and spares.

Power and energy budget: nominal, peak and endurance

For every major load, record nominal power, peak or surge power, duty cycle, protected minimum and whether the load can be deferred. Sum the loads under explicit scenarios rather than producing one daily average. The settlement should survive a defined source loss while preserving life-critical functions and a recovery path.

Primary-source bridge. NASA DRA 5.0 power studies illustrate why surface power must be treated as an architecture driver rather than a single nameplate value. [NASA NTRS — DRA 5.0 Power Requirements](#).

Energy storage needs an endurance statement. A battery may cover a protected 250 kW load for hours or only minutes depending on usable energy. The capstone should therefore show at least one timeline from source loss to load shedding, storage depletion, repair or backup generation. “Battery backup available” is not a complete statement.

Water, atmosphere and food as inventory-plus-process systems

Closed-loop percentages must be translated into makeup demand, storage and failure duration. A high recovery fraction can still require substantial resupply over months, and a processor outage can make storage the dominant variable. Track potable water, technical water, wastewater, oxygen, buffer gas if used, carbon-dioxide removal capability, food inventory and locally produced food separately.

Primary-source bridge. NASA ECLSS material provides primary operational context for atmosphere revitalisation and water recovery; the settlement budgets here remain teaching scenarios. [NASA — ECLSS](#).

The capstone should identify which stocks can be cross-used and which must never be mixed. Potable reserve should not disappear into industrial cleaning without a deliberate emergency decision. Similarly, oxygen production capacity does not replace the need for stored oxygen if a repair requires the production system itself to be depressurised or powered down.

Habitat and refuge: demonstrate independence, not just extra volume

A refuge is valuable only if it remains usable during the hazards that force people into it. Review its atmosphere, power, communications, thermal control, fire isolation, water, medical supplies and sanitation against the initiating event. A refuge electrically fed from the same

failed bus as the damaged habitat may be a room, not a refuge.

Evacuation routes should be tested with realistic obstructions, suited or injured crew, night-shift staffing and doors or hatches that may be unavailable. Timed drills reveal whether the geometric plan translates into human performance.

Maintenance and spares: design the repair organisation

Build an equipment criticality list. For each critical item, state expected failure modes, diagnostic method, line-replaceable units if any, special tools, spare location, repair time and qualification needed. Common spare interfaces are valuable because they reduce inventory variety, but commonality can also create common-cause failure if one flawed part is used everywhere.

Primary-source bridge. NASA's Systems Engineering Handbook provides the systems-engineering context for lifecycle, verification and interface discipline. [NASA — Systems Engineering Handbook](#).

The workshop needs a defined boundary of capability. It may fabricate brackets, hoses, simple seals, machined parts or electrical harnesses while remaining unable to produce high-reliability electronics, specialty bearings or pharmaceuticals. The capstone should distinguish local fabrication, local repair, cannibalisation and Earth-only replacement.

Medical and public-health readiness

Thirty residents already justify more than a first-aid kit. The settlement needs triage space, isolation capability, sterile procedures, diagnostics appropriate to foreseeable conditions, medication management, dental contingencies, injury stabilisation and a plan for prolonged care without evacuation to Earth. Medical capability must be cross-checked with power, water, refrigeration, communications and staffing.

Primary-source bridge. NASA's Human Research Program studies health and performance risks relevant to long-duration exploration. [NASA — Human Research Program](#).

Preventive health is equally important. Water quality, food hygiene, exercise, radiation exposure tracking, sleep, occupational hazards and mental workload all affect the population before an acute medical event occurs.

Mobility and rescue radius

Map the maximum distance from pressurised safety under nominal and degraded mobility. A science rover that reaches a site does not establish rescue capability. The capstone should show at least one independent recovery route, compatible towing or crew transfer where relevant, communications coverage and energy/life-support margins.

Primary-source bridge. NASA JSC rover material provides primary context for pressurised and unpressurised surface mobility concepts. [NASA JSC — Rovers](#).

Surface operations also compete for vehicles. A rover assigned to science may be the same asset needed for maintenance on a remote power field. Scheduling must protect emergency availability rather than using the entire fleet at once.

Operations centre: convert telemetry into decisions

Define which measurements are continuously trended, which alarms are immediate, which are reviewed daily and which require human interpretation. Every alarm should point to an operator action or diagnostic path. Alarm floods are themselves a hazard because they hide the initiating event. The system should prioritise alarms and preserve event chronology.

Primary-source bridge. NASA JSC spaceflight-operations material provides real mission-operations context for procedures, flight control and decision support. [NASA JSC — Spaceflight Operations](#).

Shift handover is a formal control. The outgoing team should communicate degraded modes, inhibited alarms, temporary configurations, consumables below target, open work permits and decisions awaiting evidence. The incoming team should not have to discover the settlement state by reading scattered chat messages.

Capstone assessment: evidence beats presentation quality

A polished diagram cannot compensate for an unverified dependency. During oral defence, reviewers should pick a random critical claim—“thirty residents can survive one water processor outage for four sols,” for example—and demand the complete chain: demand assumption, verified inventory, isolation boundaries, reserve policy, repair estimate, operator procedure and source or test evidence. The student passes when the reasoning remains coherent under that interrogation.

A second review should deliberately attack attractive assumptions: perfect crop harvest, full solar output, all specialists available, no maintenance overlap, instantaneous fault diagnosis. Replace them with a coupled degraded scenario and see whether the architecture still has a safe path. This is the First-Man depth test applied correctly: not more pages for their own sake, but enough technical structure that a serious candidate can be challenged without the model collapsing into slogans.

R59 commissioning board: turn a settlement design into an acceptance campaign

A thirty-person settlement should not be accepted because every subsystem has been delivered, powered once or marked green on a procurement dashboard. Acceptance is an evidence campaign. The board must show that life-critical functions remain available through representative operating states, that interfaces do not create untested common causes, and that the crew workload required to keep the system alive is itself supportable.

Build a dependency ledger with failure direction

A conventional block diagram shows that power feeds water processing and thermal control. A commissioning ledger goes further: it records how the dependency fails, how quickly the consequence appears, whether the failure is observable, what buffer exists and which recovery action is possible. A loss of electrical power may stop a pump immediately, but the human consequence can be delayed by tank inventory. A communications failure may not stop oxygen production at all, yet it can remove remote expertise during a simultaneous maintenance event. These are different dependency classes and should not receive the same generic “redundant” label.

For each critical edge, record a direction: source function → dependent function → first lost capability → time to unacceptable state. Then attach the protective mechanism: stored inventory, alternate path, manual procedure, portable equipment, isolation boundary or mission rule. The board should be able to point to the evidence that each mechanism was actually exercised.

Commission in states, not in a single ceremonial test

The acceptance campaign should include at least an uncrewed dormant state, robotic activation, minimum-care crew state, normal 30-person operations, planned maintenance, representative peak load and one or more defined degraded modes. A system that works only when every parallel unit is available has not demonstrated resilience. Conversely, forcing every subsystem into its most severe failure simultaneously can create an unrealistic test that teaches nothing. The test matrix should select credible combinations that expose interfaces and common causes.

One useful board technique is to separate **verification** from **validation**. Verification asks whether the system satisfies the written requirement: for example, whether a protected water train can deliver the specified flow. Validation asks whether the requirement and system actually support the intended mission: for example, whether the flow, storage, sampling workload and maintenance burden together let thirty people live safely through the defined resupply interval. Both are needed.

Crew workload needs a red line

Temporary manual workarounds can hide an architecture that is not operationally sustainable. If two technicians must continuously babysit a processor, the settlement may meet a flow requirement while failing the human-systems requirement. The board should therefore measure recurring preventive maintenance, fault response, sampling, inventory work, cleaning, EVA preparation and administrative coordination. Protect sleep, medical readiness, training and emergency response rather than treating every free hour as allocatable labour.

A useful stress test is to inject a life-support fault during a period already containing routine maintenance and a medical absence. If the response requires the only qualified electrical technician and the only trained medical responder at the same time, the architecture has a human common cause. Cross-training and procedure quality can reduce that exposure, but they must be demonstrated rather than assumed.

Acceptance evidence should expire when the configuration changes

Commissioning evidence belongs to a specific configuration: software version, valve lineup, filter type, battery state, crew procedure, sensor calibration and environmental condition. A major modification can invalidate part of the evidence package. Configuration control therefore

needs to answer which tests must be repeated after a change. The goal is not bureaucratic paperwork; it is to prevent a settlement from claiming yesterday's evidence for today's different system.

Board exercise — one attractive NO-GO

The cargo schedule is late and Earth wants the crew to depart during the current launch opportunity. Nominal life-support capacity supports 36 people, power supports 40 and food stocks support 45. However, sanitation N-1 support is only 28 people and the refuge has never been occupied at the full 30-person load. The board should refuse the temptation to average those numbers. The unresolved sanitation single-failure case is below planned population and the refuge evidence is incomplete. A defensible disposition is HOLD until the sanitation limitation is removed or the planned crew is reduced, and until a representative refuge test demonstrates its occupancy, atmosphere, communications and access constraints.

Primary-source bridge. NASA's Moon to Mars strategy explicitly treats exploration as an integrated system of systems and includes crew time, maintainability, reuse and interoperability as recurring architecture considerations. The acceptance framework above is a Delta-Sierra teaching method built on those systems-engineering principles, not a NASA certification procedure. [NASA — Moon to Mars Strategy and Objectives](#).

R60 capstone expansion: certify a 30-person settlement as an integrated operating system

A thirty-person settlement should not be accepted because thirty beds exist and the average power budget closes. The capstone should be defended as a chain of verified functions with explicit failure boundaries, crew workload and recovery paths. The strongest final review therefore resembles an operational readiness review: every life-critical claim is tied to evidence, every interface has an owner, every degraded state has a response, and population is admitted only when the limiting verified capacity remains above the planned occupancy with a stated margin.

Build a capacity ledger with five different meanings of “capacity”

For each critical function—water, oxygen, carbon-dioxide removal, electrical power, heat rejection, food, sanitation, medical care, shelter, communications and mobility—record installed capacity, currently available capacity, verified capacity, protected demand and N–1 capacity. Those columns prevent a recurring systems-engineering error: treating hardware that exists as hardware that can be relied upon. A spare pump in a crate does not increase current capacity until it is installed, tested and supported by the correct interfaces, tools and trained crew.

Primary-source bridge. NASA’s Systems Engineering Handbook emphasises requirements, interfaces, verification and validation across the lifecycle. The five-column capacity ledger is a Delta-Sierra teaching mechanism for making those distinctions visible in the settlement capstone. [NASA — Systems Engineering Handbook](#).

Population admission should consume only verified margin

Suppose installed oxygen capacity supports 45 people, available capacity supports 40, but the last end-to-end verification demonstrated only 34 under the current configuration. A plan to admit four new residents to a population of thirty is not justified by the installed number. It consumes the entire verified margin. The review should ask what changed since verification, whether the unverified capacity is needed for contingency, and whether the arriving crew adds capability quickly enough to offset the new demand. Population is a load on every shared subsystem simultaneously.

Commissioning has to precede occupancy

Cargo arrival is not the same as system readiness. Commissioning should include leak checks, flushing and cleanliness verification, sensor calibration, control-loop tuning, communications checkout, power quality, emergency isolation, alarm routing, spares identification and a demonstration that procedures match the as-built configuration. For life-support functions, the crew should practise the transition from nominal operation to a protected degraded mode before the function is credited to the admission gate.

Primary-source bridge. NASA’s ECLSS reference describes the coupled functions required to maintain a habitable environment. The capstone uses those functions as a reminder that life support is a network of interacting loops rather than a single machine. [NASA — Environmental](#)

Common-cause failures belong in the admission review

Two pumps do not provide useful redundancy if both depend on the same power converter, coolant inventory, software controller, contaminated fluid, inaccessible connector or maintenance procedure. For each critical “A/B” pair, identify what remains common. Then ask whether the common element can fail, how it is detected, and what alternate path exists. The review should be especially suspicious of dependencies that are physically distant in drawings but operationally common, such as one calibration source, one operator qualification or one shared network switch.

Crew time is a system resource

A settlement can pass every mass and power calculation yet be operationally impossible because maintenance, hygiene, food production, EVA preparation, science, training and administration require more qualified labour than the crew can supply. Create a weekly workload budget by skill, not just by headcount. Protect sleep, exercise, medical time, emergency reserve and training. If routine work already consumes nearly all skilled hours, the architecture has no resilience even if hardware margins look generous.

Primary-source bridge. NASA’s Human Research Program studies risks to astronaut health and performance. The workload ledger used here is a teaching application of the broader principle that human performance is part of mission capability. [NASA — Human Research Program](#).

Medical capability should be defined by functions and delay

“There is a clinic” is not a useful readiness statement. Define what can be diagnosed, stabilised and treated locally; which drugs and consumables are protected; which imaging, laboratory and dental capabilities exist; how many trained people can use them; and what happens if the primary medical officer is the casualty. Because evacuation to Earth is not an immediate option, the capstone should distinguish conditions that can be managed locally from conditions where the architecture only reduces harm and buys time.

Logistics readiness includes identification and retrieval time

Inventory mass is only part of logistics. A critical spare that cannot be located, is packaged behind inaccessible cargo, has no verified compatibility record or needs a tool that is missing is not operationally available. The review should sample retrieval tasks: select a seal, sensor, filter, cable or medical item and measure how long a crew member takes to find the correct part, confirm its configuration status and bring it to the worksite. This converts inventory management from a database claim into a demonstrated capability.

Mobility must close the rescue loop

Rovers should be assessed not only for productive range but for rescue. For each routine traverse, identify the farthest credible disabled location, the rescue vehicle, crew preparation time, travel time, casualty transfer method, communications path and energy reserve after return. A mobility system that reaches a site but cannot recover a disabled crew without violating another protected reserve is not fully commissioned.

Run an integrated casualty-and-utility exercise

A strong capstone exercise combines failures that compete for the same people. Example: a water-quality alarm occurs while one power branch is isolated for maintenance and an EVA crew reports a mobility fault outside the habitat. The exercise tests authority, prioritisation, communication, cross-training and protected resource use. The goal is not to “win” every scenario; it is to reveal where one specialist, one communication channel or one tool becomes a common cause.

Final readiness board — what evidence must be on the table

Gate	Evidence expected	Reason to HOLD
Life support	Verified production/removal capacity, alarms, isolation, consumables, degraded procedures	Capacity exists only on paper or depends on an untested common element
Power and thermal	Protected-load list, verified generation, storage, N–1 heat rejection, shedding sequence	Recovery requires consuming the emergency reserve
Human capability	Skill matrix, weekly workload, backups for critical roles, training records	One absence removes a life-critical competence
Medical	Stabilisation functions, stock coverage, diagnostics, backup provider	Essential care relies on one person or expired/unverified inventory
Logistics	Critical spares, retrieval demonstration, configuration records, next resupply window	A single credible failure has no local recovery path before resupply

The board should end with a decision that can be defended in writing: GO, HOLD pending named work, or NO-GO until the architecture changes. A capstone that always produces GO is not testing the settlement; it is only confirming its own assumptions.

R60 thirty-person defence: five boards that the capstone team must survive

Board 1 — systems interfaces

The panel selects an interface that normally stays invisible—power to a pump, cooling to an avionics cabinet, data to a valve controller, water quality between treatment stages—and asks what happens when the upstream system remains nominal but the interface fails. The student must show the physical signal used to detect the condition, the ownership boundary and the degraded procedure. This tests whether the architecture is integrated or merely a stack of subsystem chapters.

Board 2 — maintenance realism

The panel chooses a credible failure and asks for the complete repair chain: diagnosis, safe isolation, spare identification, access, tools, contamination control, functional test and restoration of configuration records. If the answer assumes a part can simply be “replaced,” the defence is incomplete. Mars maintenance includes the work around the component, and that surrounding work often dominates crew time and risk.

Board 3 — human continuity

Remove one critical person from the shift for forty-eight hours. The capstone must show which roles still have qualified coverage and which tasks are deferred. This is particularly important for medicine, electrical switching, atmosphere control, EVA leadership and specialised maintenance. A thirty-person community can have enough total labour but still be critically dependent on one individual.

Board 4 — inventory uncertainty

Declare one database quantity wrong: a critical spare is missing, a reagent lot is quarantined, or a filter inventory was double-counted. The team must identify which operational claims depended on that stock and how quickly the error propagates into capacity or recovery

planning. The lesson is that inventory accuracy is a safety function, not clerical housekeeping.

Board 5 — delayed resupply

Move the expected resupply opportunity later than planned. The team must revisit protected stocks, maintenance deferrals, medical inventory, food diversity, industrial feedstocks and spares. A capstone that is robust only when every cargo flight arrives on schedule is not demonstrating a settlement; it is demonstrating a tightly coupled expedition.

Primary-source bridge. NASA's Moon to Mars strategy emphasises objectives, architecture and sustained exploration across missions. These defence boards are Delta-Sierra teaching devices for testing whether a proposed settlement can preserve capability when assumptions are deliberately disturbed. [NASA — Moon to Mars Strategy and Objectives](#).

R60 evidence dossier: what the 30-person team should be able to hand to an independent reviewer

The final dossier should contain more than narrative design. It should include a one-page configuration summary for each critical subsystem, current revision identifiers for procedures, a list of open waivers or temporary configurations, the protected-load matrix, the latest end-to-end verification date, unresolved anomalies and the name of the person or role responsible for closure. This makes the capstone auditable by someone who did not build it.

For life support, include recent trend plots or equivalent data showing that atmosphere and water functions remain stable across realistic demand variation. For power and thermal, include the protected-load sequence and the last demonstration of degraded operation. For mobility, include rescue range and the last recovery drill. For medical capability, include stock coverage by treatment function and the backup provider path. For logistics, include critical spares whose loss would create Earth dependence longer than the settlement can tolerate.

The reviewer should also receive a map of interfaces that can propagate failure: shared power buses, data networks, coolant loops, pressure boundaries, ventilation, software services, crew qualifications and physical access routes. The map is not a complete fault tree, but it exposes where two apparently separate systems are actually coupled. A capstone that hides these couplings in different chapters is hard to challenge and easy to overestimate.

Finally, the dossier should contain a short list of explicit non-capabilities. Examples might include surgery beyond a defined level, manufacture of a particular electronic component, recovery after a certain class of habitat breach, or operation beyond a stated dust/temperature envelope. Declaring what the settlement cannot do is evidence of maturity. It prevents future planners from turning an omission into an assumed capability.

Acceptance exercise — defend one HOLD decision

The team should deliberately choose one plausible configuration in which the correct answer is HOLD. For example, all nominal capacities pass, but one critical N-1 function has only a two-percent verified margin and its repair kit is scheduled on the next cargo flight. The defence must explain why schedule pressure does not justify admission, what exact work restores the gate, what evidence will be collected, and who has authority to release the HOLD. This trains the most important capstone habit: refusing to confuse optimism with verification.

Population gates should use verified degraded capacity, not brochure or nameplate capacity. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.