

Capstone: scale a Mars settlement from 30 to 100 people



Conceptual visualization — scaling toward 100 residents changes infrastructure topology, staffing, stocks, maintenance, governance and failure consequences.

Scale power, water, habitat, food, health, maintenance, logistics and governance without simply multiplying a 30-person base.

Before starting — reconnect settlement architecture, systems engineering, logistics and the 30-person capstone before defending growth toward one hundred residents. Important quantities and assumptions are stated at first use.

Mastery objectives

- show which settlement functions scale proportionally, which are fixed and which change regime near one hundred residents
- design districts, utility trunks, health systems, logistics and staffing around failure containment and rerouting
- use staged admission gates that test N-1 capacity, maintenance backlog, qualification depth and public-health resilience
- measure autonomy by recovery delay and consequence so growth never outruns institutions or imported dependencies

1. Scaling is not multiplying by 3.33


Some consumption scales nearly with population, including parts of water and food demand. Other functions can be shared: laboratories, heavy workshops or mission control. New needs also appear: multiple work teams, separated living zones, larger rescue capacity, distributed stores and more formal governance. The design must identify these changes of regime.

2. Growth by stages

Jumping directly from 30 to 100 people makes errors expensive. A staged plan such as 30 → 60 → 100 defines admission criteria before each arrival: available power, water reserve, ECLSS capacity, refuge beds, food production, medical stocks, vehicles and qualified technicians. New crews depart only after the previous stage is demonstrated.

3. Habitat and compartmentation

A 100-person settlement should not depend on one large pressurized volume. Sleeping, work, refuge, medical, agriculture and storage are divided into isolatable zones. Connections support daily circulation but can close during fire, contamination or pressure loss. Refuge capacity must match scenarios in which an entire district becomes unavailable.



At one hundred residents, utilities and services must be sectionalised so a local fault can be isolated without orphaning the settlement. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

4. Water and atmosphere

Flow rates rise, but contamination consequences rise even more. Loops are divided into trains or zones with sampling, purge and controlled cross-feed. Large tanks buffer operations, but distributing part of the inventory prevents one incident from making the entire reserve unavailable.

5. Power and microgrids

At 100 residents a single network can propagate faults. The design uses multiple buses or microgrids that can support one another without being inseparable. Critical loads retain local autonomy. Large industrial consumers participate in shedding and are scheduled around generation, storage and maintenance.

6. Food and agriculture

Local production may cover a growing fraction of calories, but more crop area also means more light, water, nutrients, labor and heat. Imported stores remain insurance against plant disease or power loss. The project measures calories and nutrient coverage, not just harvested mass.

7. Health and population

A larger population increases the frequency of rare medical events. Skills, pharmacy, diagnostics and treatment spaces must therefore expand. Some capabilities can be shared; others need redundancy, including medical power, oxygen, sterilization and communications. Public-health management also becomes more important.

8. Maintenance and specialization

One hundred residents provide more specialists but also more equipment. Every life-critical function needs trained backup personnel. Maintenance teams manage configuration, inspections, spares and planned work while preserving emergency response capability.

Deepening: logistics for growth

The manifest carries more than consumables. It prepares the next stage with machines, structures, spares and feedstocks that raise local capacity before the next group arrives. Poor sequencing can create months in which population grows faster than infrastructure.

Deepening: governance and common services

At 100 people not every daily decision can climb to one commander. Responsibilities distribute by function and district under common safety and emergency rules. Governance handles scheduling, maintenance, resource allocation and disputes without turning every technical tradeoff into a political crisis.

9. Worked example: growth water reserve

At 3.5 L of direct potable and food-preparation water per person per day, 100 residents require 350 L/day. A 21-day reserve requires $350 \times 21 = 7,350$ L. If loss of one tank must leave at least half the reserve, two 3,675 L tanks do not preserve 21 days after one tank is lost. More total volume or another source is required. Redundancy sizing is therefore more than nominal total capacity.

Calculated case study: scaling from 30 to 100 people without simply multiplying by three

TEACHING ASSUMPTION — Infrastructure has 60 kW of fixed loads independent of population and 4 kW of variable loads per person. Add a 30% contingency margin at the 100-person stage.

Let P_f be fixed load in kW; p variable load in kW/person; N population; $P(N)$ power before margin; and r margin, dimensionless.

At 30 people: $P(30) = P_f + p \times 30 = 60 + 4 \times 30 = 180$ kW. At 100 people: $P(100) = 60 + 4 \times 100 = 460$ kW. With a 30% margin: $P_{\text{plan}} = 460 \times 1.30 = 598$ kW.

Growth is not a simple factor of $100 \div 30$ because part of the infrastructure is fixed and part varies with population. The distinction forces the design to identify what can be shared before sizing the expansion.

10. Integrated exercise

A 60-person base uses 900 kW average and plans to add 40 residents. Each new resident adds 4 kW of direct and indirect load in this design stage. Calculate the new load and margin if available generation remains 1.15 MW, then explain why the power calculation alone cannot authorize the arrival.

11. Reasoned solution

Forty new residents add $40 \times 4 = 160$ kW, taking load to 1,060 kW. Margin on 1,150 kW is 90 kW, about 7.8%. That may be inadequate under maintenance, dust, aging or peaks. Water, ECLSS, refuge beds, food, health, mobility and workforce also need to pass their own gates.

12. Final deliverable

Produce the $30 \rightarrow 60 \rightarrow 100$ growth plan with zone architecture, capacities at each stage, go/no-go criteria, mass/power/water/food budgets, medical strategy, maintenance organization, governance and ten remaining Earth dependencies. The design must also show what happens if one growth cargo is lost.

Final capstone: turn a thirty-person outpost into a hundred-person settlement

Scaling from thirty to one hundred residents is a change of regime, not a copy-and-paste exercise. Some needs rise almost linearly with population, some infrastructure can be shared, and entirely new functions become necessary once the settlement has several work teams, multiple living districts and a larger number of simultaneous medical, maintenance and logistics demands. The capstone therefore asks the learner to design a growth pathway rather than a single final snapshot.

The central rule is **capacity before population**. New residents should arrive only after the previous stage demonstrates the power, thermal rejection, water, atmosphere, food, refuge, medical, mobility and maintenance capacities required for the next stage. Cargo that expands

those capacities must therefore arrive and be commissioned before the people who depend on it.

The exercise uses planning values to teach systems integration. These values are not statements that NASA, ESA or another agency has selected this exact settlement architecture. Every numerical assumption in the final dossier should be labeled as sourced data, derived value or teaching/planning assumption.

1. Population trajectory and admission gates

A useful growth plan defines stages such as $30 \rightarrow 45 \rightarrow 60 \rightarrow 80 \rightarrow 100$ rather than one jump. Each stage has entry criteria, a commissioning period and an observation period. Population growth is delayed when one safety-critical gate is not met, even if other systems have excess capacity.

Admission gates can include available power after N-1 failure, thermal-rejection margin, protected potable-water reserve, atmosphere-control throughput, refuge places, medical beds and oxygen, food coverage, mission-ready rovers, qualified maintenance staff and validated communications capacity. The exact thresholds require mission design, but the principle is universal: growth should be evidence-based.

Human skills are part of capacity. A population of one hundred may justify specialization, but critical functions still need cross-trained backups. The settlement should avoid situations where only one person can perform a life-critical maintenance action or interpret a diagnostic instrument.

2. Power and microgrids at the hundred-person scale

Electrical demand has both fixed and population-dependent components. Mission control, network infrastructure, large workshops and some industrial systems may not scale directly with headcount. Habitat life support, food preparation, personal equipment and many service loads do. The design should separate these categories so expansion does not blindly multiply every number by 100/30.

$$P_{\text{total}}(N) = P_{\text{fixed}} + N \times p_{\text{person}} + P_{\text{industry}}(N)$$

If fixed infrastructure is 70 kW, population-dependent demand is 3.2 kW per person and scheduled industry averages 180 kW at the 100-person stage, average modeled demand is $70 + 100 \times 3.2 + 180 = 570$ kW. Peak and contingency demand must be analyzed separately.

At larger scale, the electrical architecture should prevent one fault from blacking out every district. Multiple buses or microgrids can support each other while retaining islanding capability. Critical loads such as refuge, medical care and communications should retain local energy or alternate supply so cross-settlement faults do not become immediate survival events.

Industrial loads become useful flexibility. Oxygen production, some ISRU operations, machine tools, thermal storage and some agricultural lighting can shift in time when power is constrained. The schedule should therefore integrate generation forecast, storage state and production deadlines rather than treating every load as continuous.

3. Thermal rejection becomes a growth gate

Adding electrical power without adding heat-rejection capacity can make the new generation unusable. Nearly all consumed electrical energy eventually becomes heat inside or near the settlement. Greenhouses, computing, workshops and chemical plants add thermal loads that may peak at different times.

$$Q_{\text{reject}} = Q_{\text{internal}} + Q_{\text{solar}} + Q_{\text{process}} - Q_{\text{useful_recovered}}$$

The settlement should account for heat recovered for water warming, habitat heating or process preheat before sizing final rejection. But heat recovery cannot eliminate the need for rejection when the receiving loads are already warm or when seasonal conditions change.

Thermal capacity should be included in every population admission gate. A stage can have enough electricity and water yet still be unsafe because radiator, loop or heat-exchanger capacity has no margin after one component is unavailable.

4. Water, sanitation and distributed storage

At one hundred residents, daily flow is larger and contamination affects more people. The architecture benefits from multiple treatment trains, isolatable storage and sampling points that allow one zone to be quarantined without losing the entire inventory. A single very large tank can be efficient but creates a concentration of consequence.

$$V_{\text{reserve}} = N \times q_{\text{emergency}} \times d \times (1 + \text{margin})$$

If emergency planning uses 9 L/person/day inside a defined accounting boundary, one hundred residents, 10 days and 25% margin, reserve is $100 \times 9 \times 10 \times 1.25 = 11,250$ L, or 11.25 m³.

This is a teaching example; actual emergency water needs depend on system architecture and conservation policy.

Sanitation capacity must scale with flow and with the ability to isolate contamination.

Wastewater processing, hygiene, microbial monitoring and solid-waste interfaces should be treated as settlement infrastructure, not secondary services added after habitat expansion.

5. Atmosphere control and compartmentation

One hundred people should not depend on one undivided pressurized volume. Living, medical, laboratory, industrial, storage and agricultural zones need controlled interfaces so fire, contamination or pressure loss can be contained. Connectivity remains important for daily life, but the settlement should be able to close boundaries without trapping people away from refuge.

Atmosphere-control capacity scales with metabolic load and with leakage, process releases and contingency cases. The capstone should show normal processing, N-1 capacity, emergency oxygen reserve, carbon-dioxide removal, humidity control and the isolation logic between districts.

Refuge capacity should be scenario-based. If one district becomes unavailable, the remaining refuge spaces must support the displaced population for the time needed to repair or evacuate the affected zone.

6. Food, agriculture and inventory resilience

A larger population can justify more local food production, but local production creates additional dependencies: lighting, irrigation, nutrients, seed stock, labor, pest or microbial control, processing, cold storage and thermal rejection. The settlement should track edible calories and key nutrients after losses rather than crop mass alone.

$$\text{Food_coverage_days} = \text{Energy_usable} / (\text{N} \times \text{kcal_plan})$$

If usable stored food contains 27,000,000 kcal and the planning assumption is 3,000 kcal/person/day for 100 residents, simple coverage is $27,000,000 / 300,000 = 90$ days. Nutrient balance, shelf life and food safety must be checked separately.

Imported food remains strategic insurance even when local production becomes substantial. A plant disease, power shortage or greenhouse contamination can reduce output suddenly. Protected reserve should be sized against credible production-loss scenarios rather than average harvest.

7. Habitat becomes urban infrastructure

At one hundred residents, habitat planning begins to resemble a small town. Sleeping, hygiene, food, recreation, medical care, education, operations, workshops and storage compete for space and quiet. Movement paths need to support daily circulation and emergency isolation. The design should provide private or semi-private retreat as well as shared social spaces because chronic crowding can become a performance and health issue.

Expansion modules should connect through standardized interfaces for pressure, power, data, fluids and emergency closure. Standardization reduces commissioning time and spare variety. The final layout should show how a new district can be added without taking the entire settlement offline.

Fire zones and refuges should remain distributed. A central refuge that requires every resident to cross the same corridor can create a single evacuation bottleneck.

8. Industry, maintenance and qualification

A hundred-person settlement owns more equipment and can support deeper specialization. The workshop may expand from repair toward local production of structural parts, seals, pipes, simple electronics assemblies or chemical products. However, manufacturing a part is not the same as qualifying it for a life-critical use.

The settlement should maintain a **make / repair / refurbish / stock / import** matrix for critical items. Each local-production pathway lists required feedstock, machines, tooling, measurement equipment, software, skills and quality evidence. This exposes hidden Earth dependencies such as a cutting tool, sensor, catalyst or calibration artifact.

Maintenance staffing also needs queue analysis. If preventive and corrective maintenance consumes nearly all technician time during normal operations, the settlement has no capacity left for anomalies or expansion work.

$$W_{\text{maintenance}} = H_{\text{required}} / H_{\text{available}}$$

If planned maintenance needs 1,200 person-hours per month and the trained team can provide 1,600 effective hours after rest, training and emergency reserve, workload ratio is 0.75. The unused 25% is not waste; it is margin for corrective work and growth tasks.

9. Mobility and surface logistics

A larger settlement creates more simultaneous routes between habitats, landing zones, power fields, mines, science sites and construction areas. Fleet design should separate crew-rated transport, cargo movement, rescue and specialized industrial vehicles where practical. Counting vehicles is not enough; the plan needs mission-ready availability and charging capacity.

Surface roads, graded corridors, navigation aids, charging points and shelters become infrastructure. Improvements that appear expensive for thirty residents may become economical when used by hundreds of trips per year.

Rescue geometry should be redrawn at each growth stage. More residents and more simultaneous field teams can exceed the capacity of one rescue rover even if total fleet size has grown.

10. Health system scaling and public health

As population increases, the expected frequency of uncommon medical events rises. The settlement can justify additional diagnostic equipment, isolation capability, pharmacy stock and clinical specialization. It also has to plan around simultaneous cases: one surgical event should not eliminate routine medical capacity for everyone else.

Public-health functions become more important. Food safety, water surveillance, infection control, occupational exposure, mental health and injury prevention can affect many residents at once. Medical planning should therefore include prevention and population surveillance, not only treatment rooms.

Critical clinical functions need independent power, oxygen, communications and sterile-supply pathways. The health system should appear explicitly in microgrid and logistics diagrams rather than as a box labeled “hospital.”

11. Governance and specialization at one hundred residents

A population of one hundred cannot route every technical decision through one commander. Authority should distribute to functional and district leads under common emergency rules. This is not bureaucracy for its own sake; distributed authority reduces bottlenecks and lets specialists manage routine operations while preserving coordinated crisis command.

Resource allocation becomes more visible because multiple teams may compete for workshop time, vehicle access, laboratory capacity or scarce imports. Scheduling systems and decision logs should make the basis of these choices transparent. The settlement should distinguish safety priorities from social or scientific priorities so technical limits are not hidden inside general debates.

Growth also creates training obligations. New residents need orientation to emergency states, refuge routes, communications procedures, maintenance reporting and local governance. A system that works only because the original thirty founders remember unwritten rules will fail as population grows.

12. Earth dependence and progressive autonomy

Autonomy should be measured by dependency, not rhetoric. The final dossier identifies which functions still require Earth-supplied hardware, software, consumables, expertise or raw materials. Some dependencies may remain rational for decades because local production would be too complex or inefficient.

A useful autonomy roadmap ranks dependencies by consequence and replacement horizon. Losing an imported decorative item is irrelevant; losing a unique membrane, drug, bearing, catalyst or computing component may be mission-critical. Local industry should therefore target high-consequence dependencies first.

Communications delay does not eliminate Earth collaboration. Earth can continue to provide science, design support, software, medical consultation and logistics planning. Progressive autonomy means the settlement can remain safe and productive when those links are slow or temporarily unavailable.

Master capacity calculations

Population-dependent demand

$$D(N) = D_{\text{fixed}} + N \times d_{\text{person}}$$

Use this form only when a fixed component and a per-person component are physically meaningful. For example, if a communications center requires 20 kW fixed and user terminals plus support loads add 0.15 kW/person, the 100-person modeled load is 35 kW.

N-1 capacity check

$$\text{Margin}_{N-1} = \text{Capacity_after_largest_single_loss} - \text{Essential_demand}$$

If three identical 300 kW power blocks provide 900 kW and loss of one leaves 600 kW while essential demand is 520 kW, N-1 margin is 80 kW. This does not address common-cause failures that could remove more than one block.

Refuge occupancy margin

$$M_{\text{refuge}} = \text{Beds_or_places_available} - \text{Population_displaced_in_design_case}$$

If the largest isolatable district holds 28 people and alternate refuges together provide 35 places, occupancy margin is seven. Air, water, thermal and sanitary capacities must also support those occupants.

Inventory coverage

$$\text{Coverage} = \text{Stock_usable} / \text{Consumption_rate}$$

The consumption rate should reflect the contingency being analyzed. Routine average use can overstate resilience if a failure increases demand.

Growth readiness index

$$\text{Readiness} = \min(\text{gate_power}, \text{gate_thermal}, \text{gate_water}, \text{gate_ECLSS}, \text{gate_food}, \text{gate_medical}, \text{gate_refuge}, \text{gate_mobility}, \text{gate_staffing})$$

Each gate can be represented as 1 when its approved requirement is met and 0 when it is not. Taking the minimum makes the logic explicit: one failed critical gate blocks population growth. This is a governance representation, not a substitute for engineering margins.

Reference growth case: 30 → 60 → 100 over 1,000 sols

Stage A — stabilize thirty residents

The first stage focuses on demonstrating closed-loop performance and maintenance rather than maximizing population. The settlement records real power variability, water losses, crop yields, medical consumption, rover energy and technician workload across seasonal conditions. Critical spares are adjusted from actual failure data.

Before growing, the settlement proves that one major habitat compartment can be isolated, one main power element can be unavailable without loss of refuge, and a delayed cargo does not immediately threaten survival.

Stage B — predeploy capacity for sixty residents

Cargo adds habitat volume, a second or expanded atmosphere train, water storage, additional microgrid capacity, thermal rejection, food-production equipment, medical stocks and maintenance capability. These systems are commissioned robotically or by the existing crew before the new residents depart.

Staffing plans identify which new specialists solve current bottlenecks. The arrivals should increase capability rather than only increase consumption. Training records show backup competence for life-critical functions.

Stage C — operate sixty residents through a proving period

The settlement pauses growth long enough to observe whether the larger population reveals new bottlenecks. Waste handling, meal preparation, hygiene, medical scheduling, communications, workshop queues and refuge drills can expose limits that were invisible in a thirty-person system.

A growth review compares measured margins against the 100-person design. Infrastructure is redesigned where scaling assumptions were wrong rather than simply adding another identical module.

Stage D — predeploy the hundred-person network

The next cargo phase adds distributed capacity: more than one power island, additional storage, secondary medical space, expanded workshop and industrial capability, surface logistics nodes and habitat zones that can isolate from one another. Growth is postponed if thermal or maintenance capacity lags behind headline power generation.

Stage E — admit the final population only after gate closure

The one-hundred-person stage opens only when every critical gate is demonstrated. The final review includes failure cases, not merely nominal capacity. It asks what remains after the largest credible single loss and after selected common-cause scenarios.

Six injected events for the capstone

Event 1 — population arrives faster than habitat

A transport schedule advances while one habitat module is delayed. The learner must decide whether temporary crowding remains inside refuge, hygiene, privacy and life-support limits. The correct answer may be to delay arrivals rather than “make room” by consuming contingency capacity.

Event 2 — thermal rejection becomes the limiting capacity

New solar arrays raise electrical generation, but radiator expansion is late. The learner recalculates which industrial and agricultural loads can operate without exceeding thermal limits and identifies whether heat recovery or scheduling can defer some rejection demand.

Event 3 — a critical imported feedstock is delayed

A consumable for a local chemical process will arrive one launch window later than planned. The learner calculates coverage, alternative uses, substitution possibilities and the consequences for downstream systems. The exercise should reveal whether “local production” was still dependent on an imported enabling material.

Event 4 — the workshop cannot qualify a vital part

The settlement can manufacture the geometry of a failed component but lacks the test capability to prove its material or performance. The learner must separate emergency use under explicit risk from routine qualification, and decide what test equipment or design changes reduce this dependency in future.

Event 5 — the health system is saturated

A multi-casualty event occupies the main clinical area while routine needs continue. The learner checks alternate treatment space, oxygen, power, sterile supplies, staffing and communication support. The event tests whether the larger settlement gained real medical capacity or only a larger patient population.

Event 6 — growth cargo is lost

One cargo vehicle carrying habitat, spare parts and a microgrid element is lost before arrival. The learner must identify which next-stage admissions are blocked, which activities can continue and whether any critical dependency was concentrated in that cargo. The manifest should have been designed to avoid one loss destroying all capacity of a single kind.

Master evidence table required in the final dossier

Domain	30 people	60 people	100 people	N-1 / contingency evidence	Earth dependency
Power	Measured baseline	Expanded microgrid	Distributed buses	Largest single loss + peaks	Converters, cells, controls
Thermal	Seasonal rejection	Additional loops	District capacity	Loop/radiator unavailable	Pumps, coatings, sensors
Water	Recovery + storage	Second train	Distributed tanks	Contamination/isolation case	Membranes, sorbents
Food	Imported + crops	Expanded production	Diverse production	Greenhouse outage	Seeds, nutrients, equipment
Medical	Core capability	More diagnostics	Distributed capability	Multi-casualty scenario	Drugs, sterile consumables
Mobility	Small fleet	Separated roles	Network logistics	Rescue while one rover down	Cells, tires/wheels, electronics
Maintenance	Generalists	Growing specialization	Dedicated teams	Surge workload	Tooling, metrology, spares

The numbers in this table belong to the learner's chosen architecture. The table forces every stage to show both nominal capacity and degraded evidence. A cell labeled “adequate” without a value, method or test is not sufficient.

Progressive exercises with solutions

- ▶ Exercise 1 — Fixed and variable power
 - ▶ Exercise 2 — Water reserve
 - ▶ Exercise 3 — Food coverage
 - ▶ Exercise 4 — N-1 power
 - ▶ Exercise 5 — Maintenance workload
 - ▶ Exercise 6 — Growth gate
-

Final design package

The final submission should include a 1,000-sol growth timeline, population stages, cargo sequence, commissioning plan, settlement zoning, power and thermal budgets, water and atmosphere balances, food strategy, medical capability, fleet and rescue map, maintenance workload, local-industry matrix, governance structure, Earth-dependency register, protected reserves and explicit go/no-go criteria before each population increase.

For each of the six injected events, show the initial state, information available, decision authority, immediate action, resource impact, recovery path and change to the architecture. At least one event should force a delay in population growth; otherwise the design has probably made its gates too weak.

The strongest capstone does not claim perfect autonomy. It shows which dependencies remain, why they remain, how long the settlement can tolerate interruption and which next investment reduces the largest consequence. That is the difference between a visionary image and an architecture that can be interrogated, tested and improved.

Interactive beginner glossary

- admission gate — go/no-go condition for growth.
- N-1 capacity — capacity after one major element is lost.
- microgrid — controllable local power network.

- commissioning — verification before service.
- Earth-dependency register — list of remaining external dependencies.
- workload ratio — required work divided by available work capacity.
- compartment — independently isolatable settlement zone.
- contingency reserve — resource held back for adverse events.

Scaling from thirty to one hundred residents changes the architecture, not only the numbers

The final capstone should challenge the assumption that a settlement is a small base multiplied several times. Population growth changes topology, staffing, traffic, maintenance, social organization, inventories and the consequences of outages. Some systems gain efficiency from scale; others require new distribution layers and new isolation boundaries.

Fixed, proportional and threshold loads

Classify each demand before scaling it. A central communications mast may be largely fixed. Food, hygiene and sleeping volume have strong population-dependent components. A second clinic, additional transformer, new water-treatment train or another pressure zone may appear only after a threshold is crossed. The correct model can therefore be piecewise rather than linear.

For every subsystem, the student should state the scaling law used and why. If the law is uncertain, calculate at least a low and high case rather than hiding uncertainty behind a single precise number.

Distribution networks become first-class systems

At thirty people, one compact habitat cluster may tolerate short internal routes. At one hundred, power cables, water pipes, data links, ventilation, corridors and emergency egress can become networks with their own failure modes. Isolation valves, sectionalizing switches and local buffers can prevent one fault from disabling the whole settlement, but they also add components and maintenance.

Sectional capacity margin

$$M_{\text{section}} = (C_{\text{section}} - D_{\text{peak}}) / D_{\text{peak}}$$

If a distribution section can deliver 180 kW and the modeled local peak is 150 kW, section margin is $(180-150)/150 = 0.20$, or 20%. That margin must be tested after credible rerouting. If an adjacent section fails and 50 kW is transferred into this one, the same hardware may become overloaded.

Maintenance organization must scale before equipment count does

A larger settlement requires planning, configuration control, stores management, calibration, inspection scheduling and technical records. Informal knowledge held by a few founders becomes fragile. The capstone should therefore include roles for maintenance planning, spare-part control, safety, quality and training, while avoiding unnecessary bureaucracy that consumes people without improving reliability.

Skill coverage can be mapped by function. For every life-critical discipline, identify primary competence, backup competence and the time needed to restore capability if both are unavailable. A hundred residents offer more potential specialists than thirty, but also create more equipment and more simultaneous work.

Stocks should be expressed as both quantity and recovery time

A larger settlement may hold more absolute inventory while having less buffer time because consumption and industrial demand rise. The capstone should convert oxygen, water, food, medicines, filters, lubricants, gases and critical spares into time under nominal and degraded conditions. It should also state which stocks can be produced locally and which still depend on Earth.

Growth creates construction interference

Expansion work can threaten existing operations. Excavation may damage buried services, new electrical loads may change protection coordination, construction dust may enter airlocks, and commissioning a new process can alter shared utilities. Growth plans therefore need simultaneous-operations control just like maintenance work.

Governance must distinguish settlement-wide and local decisions

At one hundred people, not every choice should reach one commander or council. Local technical teams need delegated envelopes, while settlement-wide policies remain accountable through a broader process. Emergency command succession, technical dissent and protected safety limits must still work when several incidents occur at once.

Final integrated scenario: growth during degraded operation

Imagine the settlement is commissioning a new greenhouse and industrial bay when one power section develops a fault. At the same time, a dust event reduces external operational flexibility and a medical stock shipment is delayed in the teaching scenario. The student must decide which commissioning work stops, how loads are rerouted, whether the older greenhouse receives protected priority, which inventories are rationed, what maintenance staff are reassigned and which expansion milestone is formally delayed.

The response should include power and water balances, time-to-consequence for at least five functions, a staffing table, a configuration map, a decision log and explicit restart criteria. It should also identify one decision that looks economically attractive but weakens resilience, and one apparently inefficient reserve that is justified because it protects recovery.

Graduation standard for the capstone

The student should be able to defend the architecture against a skeptical review panel. Every important number must have an assumption, unit and source or be clearly labeled as a teaching scenario. Every critical function must have a failure path and recovery plan. Every major visual should correspond to the real architecture rather than decorate it. The final design is not “the” correct Mars settlement; it is a transparent, testable argument whose consequences can be recomputed when assumptions change.

Hundred-person capstone laboratory: scaling exposes fixed loads, variable loads and new failure domains

A settlement does not scale by multiplying every 30-person number by 100/30. Some loads are fixed, some grow with residents, and others appear only when new facilities or redundancy tiers are introduced.

First-pass mixed power model

$$P_{\text{total}} = P_{\text{fixed}} + N \times p_{\text{person}}$$

Teaching scenario. Assume 80 kW of fixed infrastructure and 1.2 kW average variable demand per resident. At $N = 100$: $80 + 100 \times 1.2 = 200$ kW. Applying a 20% planning margin gives $200 \times 1.20 = 240$ kW.

Limit. This is an intentionally simple model. Workshops, greenhouses, batch chemical plants, charging stations and emergency loads can create peaks unrelated to headcount.

Consumable buffer time

$$D_{\text{buffer}} = m_{\text{usable}} / \dot{m}_{\text{net}}$$

If a protected consumable stock is 20,000 kg and net settlement use after local recovery/production is 450 kg/day, ideal buffer time is $20,000/450 = 44.4$ days. If a process failure doubles net use to 900 kg/day, the same stock lasts only 22.2 days.

Interpretation. Scaling should therefore be tested in degraded states, not only nominal averages.

DISTRICTS

sleep · work · refuge

**REDUNDANT
TRUNKS**

power · water · data

**COMMON
SERVICES**health · logistics ·
industry**GOVERNANCE**ownership · priorities ·
recovery

Teaching composite: at 100 residents, a settlement becomes a network of districts and services with redundant trunks and explicit maintenance ownership.

Settlement-scale studio: when one hundred residents turn equipment into infrastructure

At one hundred residents, the architecture changes category. A compact outpost can rely on proximity, informal coordination and a small number of central utility rooms. A larger settlement needs distribution networks, sectional isolation, multiple refuge zones, explicit public-health capability, scheduled maintenance planning and governance that can make local decisions without losing settlement-wide safety constraints. The capstone should therefore be reviewed as a networked settlement rather than as a thirty-person design multiplied by 3.33.

Network capacity must survive rerouting

Power, water, data and ventilation should be divided into sections that can isolate faults. But isolation moves demand. A power feeder that is comfortable in nominal operation can overload when an adjacent section is lost and critical loads are transferred onto it.

N-1 sectional margin

$$M_{N-1} = (C_{\text{remaining}} - D_{\text{critical}}) / D_{\text{critical}}$$

Question. After the loss of one defined element, how much remaining capacity exists above protected demand?

Example. Three 180 kW feeders serve a network. A fault removes one feeder. If the two remaining feeders can each provide 180 kW, remaining capacity is 360 kW. If protected demand after load shedding is 300 kW, margin is $(360-300)/300 = 0.20$, or 20%.

Limit. The arithmetic assumes the physical network can actually route 300 kW through the remaining switches and cables. Topology and component ratings must be checked, not merely generation capacity.

Public health appears before a hospital does

With a larger population, communicable disease, foodborne illness, water contamination and occupational exposure can affect several people at once. The medical system therefore gains surveillance and prevention functions: symptom reporting, isolation capacity, environmental sampling, vaccination or prophylaxis policy where relevant, food and water traceability, sanitation monitoring and the ability to investigate clusters. This is infrastructure because it protects population capacity, not merely individual treatment.

Specialization creates both capability and dependency

A hundred residents allow dedicated electricians, medical staff, process engineers, greenhouse specialists and maintenance planners. That improves expertise but creates a new risk: critical knowledge becomes concentrated. The settlement should map each essential role to qualified backups, training pathways, documentation and cross-discipline handover.

Critical-skill coverage ratio

$$R_{\text{skill}} = N_{\text{qualified}} / N_{\text{minimum}}$$

If a function needs at least two qualified people to cover shifts and emergencies, but only three residents are qualified, $R_{\text{skill}} = 3/2 = 1.5$. That looks comfortable until one person is injured and another is committed to a rescue task. The capstone should therefore test simultaneous unavailability, not only headcount.

Industrial ecology should close useful loops selectively

Local industry does not mean manufacturing everything. The settlement should identify high-leverage loops where local processing reduces imported mass or shortens recovery time: water, oxygen, selected gases, simple structural parts, seals, filters, cleaning reagents, regolith products and repair stock. More complex electronics or medical products may remain Earth-dependent for a long time. The capstone should distinguish local production capability from merely having raw material.

Growth gates should be reversible

Population admission is the most consequential growth decision. Before a new cohort arrives, the settlement should demonstrate spare capacity in habitat, power, thermal rejection, water, atmosphere, food, medical care, refuge, logistics and maintenance labour. If one gate is weak, the correct action may be to delay population growth while still installing infrastructure. Growth should not consume the very margin needed to recover from the next failure.

Population support margin

$$M_{\text{pop}} = (N_{\text{supported}} - N_{\text{present}}) / N_{\text{present}}$$

If the demonstrated protected configuration supports 120 residents while 100 are present, margin is $(120-100)/100 = 0.20$, or 20%. The number must be computed separately for each critical resource. The settlement is constrained by the smallest credible margin, not the average.

Final capstone evidence package

The learner should submit one dependency map, one utility topology, one protected-load table, one population growth sequence, one maintenance staffing model, one medical/public-health capability map, one industrial dependency matrix, one emergency refuge plan and one recovery scenario. Every number must be traceable to an assumption, an official or scientific source, a previous Academy calculation, or an explicitly labeled teaching scenario.

Graduation review — refuse one attractive expansion

A new greenhouse would increase food autonomy but consumes 55 kW continuous power and 18 kW additional average thermal-rejection capacity. The settlement currently has 25% nominal power margin but only 12% N-1 thermal margin. Decide whether to commission the greenhouse immediately, operate it at partial load, or delay it. State what additional evidence or hardware would change your decision.

► Reasoned solution

First-Man graduation project: one hundred people is a different kind of system

Growing from thirty to one hundred residents is not achieved by multiplying every tank, greenhouse and habitat by 3.33. At larger population, the settlement develops districts, trunks, specialist teams, public-health functions, shift work, internal logistics, governance layers and maintenance backlogs. Some systems gain efficiency through scale; others become more dangerous because a single shared trunk or software service now affects many more people. The architecture must therefore change topology as well as capacity.

The capstone should treat growth as a sequence of reversible gates. Each new population increment is accepted only after the settlement demonstrates the capacity, resilience, skills, stocks and evacuation/refuge arrangements needed for that increment. Population is then a controlled load on the infrastructure rather than an irreversible promise made before evidence exists.

Population support margin should use verified, not nominal, capacity

$$N_{\text{supported}} = \min_j (C_{j,\text{verified}} / q_{j,\text{person}})$$

Starting question

When several life-support and logistics functions have different capacities, what sets the maximum population that the verified system can support?

Read aloud

Read: "supported population is the minimum, across critical functions j , of verified capacity divided by per-person requirement."

Symbols, pronunciation and meaning

$C_{j,\text{verified}}$ is demonstrated capacity of critical function j ;
 $q_{j,\text{person}}$ is the corresponding per-person demand under the design case; the minimum operator selects the most restrictive function.

Units

If capacity is L/day and per-person demand is L/person/day, the ratio gives people. The same dimensional cancellation applies separately to oxygen, food energy, sanitation throughput or another capacity expressed per day.

Origin and status of values

Verified capacities come from commissioning and operational evidence. Per-person demand comes from mission assumptions and measured settlement behaviour. Each ratio must use consistent boundaries and reserve rules.

Why this operation

Each critical function independently sets a population ceiling. The settlement can support only the smallest of those ceilings because surplus capacity in one function cannot replace a deficit in another.

Substitution and calculation

Teaching case: oxygen system supports 118 people, water 126, food logistics 104 and sanitation 112 under the chosen reserve policy. $N_{\text{supported}} = \min(118, 126, 104, 112) = 104$ people.

Calculator entry

Compute each capacity/demand ratio separately, then select the smallest. Do not average them.

Mental estimate

The food value is visibly the smallest, so it controls the result before any precise arithmetic is needed.

Independent check

At 104 people, verify every other function still has capacity at or above its corresponding demand. If not, the ratios were built from inconsistent assumptions.

Physical or operational interpretation

With a target of 100 residents, only four-person headroom remains in the limiting food/logistics function in this teaching case. That may be too little for a final design margin.

Plain-English translation

The colony is only as populous as its tightest critical system allows.

Variation / sensitivity

If food-logistics capacity improves from 104 to 120 people, sanitation at 112 becomes the next bottleneck. Growth engineering therefore moves bottlenecks rather than eliminating the concept of a bottleneck.

Limit / assumption

This screening equation treats each capacity as a scalar and does not capture peak loads, storage autonomy, failure states, distribution topology or crew skill.

What this does not prove

Supporting one hundred people nominally does not prove that one hundred can survive a major branch failure or a long resupply delay.

Boundary case to test

If one critical function is not commissioned, its verified capacity should not be replaced by installed nameplate capacity. The growth gate remains closed until evidence exists.

At one hundred people, distribution networks become first-class systems

Generation

power, water, oxygen,
food, data

Redundant trunks

isolation, cross-ties,
metering

Districts

habitation, industry,
science, refuge

Local recovery

storage, bypass, repair,
shelter

Distribution must allow a damaged district to be isolated without shutting down the whole settlement. Cross-ties can provide alternate supply, but they also create paths for faults or contamination to spread. The network therefore needs valves, breakers, data segmentation,

measurements and procedures that make reconfiguration deliberate rather than improvised.

Specialisation increases capability and creates human single points of failure

One hundred residents allow deeper expertise: surgery, electrical protection, crop pathology, pressure-vessel inspection, software security, analytical chemistry and other functions can be staffed more professionally than in a tiny outpost. But specialisation can also make knowledge disappear into individuals. For each critical discipline, build at least a primary and backup capability, preserve procedures and training records, and avoid scheduling all qualified people into the same high-risk activity.

Population-health surveillance at one hundred residents

A larger settlement needs surveillance of respiratory illness, gastrointestinal disease, food safety, water quality, sanitation and mental-health stressors across groups, not only treatment of individual patients. Isolation space, contact reduction, ventilation control and continuity of essential work become system variables. Medical architecture therefore expands from a clinic into a population-health function.

Industry should close the loops that buy the most autonomy

Trying to manufacture everything locally can consume more equipment, skills and energy than it saves. Prioritise products by criticality, resupply burden, shelf life, production feasibility and the consequence of shortage. Simple high-mass or frequently consumed materials can be strong local-production candidates; complex electronics may remain import-dependent while the settlement develops testing, repair and cannibalisation capability.

Governance must scale before conflict scales

At thirty people, everyone may know every operational detail. At one hundred, information and authority need structure: duty teams, technical authorities, community decisions, incident command, independent review and transparent resource accounting. The goal is not bureaucracy for its own sake. It is to prevent a larger population from depending on informal knowledge that only a few founders carry in their heads.

Graduation scenario: refuse growth when only the average looks good

The settlement has average power, water and food capacity for 110 people and proposes to admit twenty new residents, rising from 85 to 105. However, the western habitation district has only one water trunk, the second medical shift lacks an independent advanced-care operator, and a critical food-processing spare has a twelve-month resupply lead time with only one unit in stock. The correct graduation answer is to delay or stage growth until those specific bottlenecks are closed. An average capacity above 105 is not a substitute for topology, staffing and spares.

Population growth is a reversible evidence programme, not a single target date. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Hundred-person settlement board: growth becomes a controlled infrastructure programme

At one hundred residents, Mars settlement design changes from “base systems plus more capacity” into an infrastructure problem. Distribution networks, district isolation, public health, shift coverage, specialist training, internal logistics, construction interference and governance become first-order design variables. The graduation project should therefore prove that growth can pause safely at every stage and that the settlement never depends on a future expansion to rescue an already under-capacity present.

Use staged population gates, not a single target date

Define at least three operating plateaus—for example thirty, sixty and one hundred residents. Before each increase, verify water, atmosphere, food logistics, power, thermal rejection, refuge, sanitation, medical capacity, maintenance staffing, mobility and critical skill coverage. The next group arrives only when the previous plateau has accumulated enough operating evidence to expose hidden maintenance and workload burdens.

Primary-source bridge. NASA’s Moon to Mars strategy describes a gradual build-up approach for exploration systems and operations; the population plateaus here are Delta-Sierra teaching gates. [NASA — Moon to Mars Strategy and Objectives](#).

This avoids a common planning error: predeploying hardware whose nominal capacity appears adequate, then discovering after growth that maintenance, spares, distribution or staffing cannot support the same number of people. Capacity is not only hardware throughput; it includes the organisation required to keep that throughput available.

Settlement admission headroom using the limiting verified function

$$H_{\text{admission}} = N_{\text{supported,limiting}} - N_{\text{present}}$$

1 – Concrete question

How many additional residents could the currently verified limiting function support before reaching its demonstrated population ceiling?

2 – Intuition

Find the weakest verified population capacity, then subtract the people already present.

3 – Quantities

Use the supported-population result from each mandatory function and select the minimum; record current population separately.

4 – Formula

Admission headroom equals limiting supported population minus present population.

5 – Read aloud

“H admission equals limiting supported population minus present population.”

6 – Symbols

H is a count of people; $N_{\text{supported,limiting}}$ is the smallest verified population capacity across required functions; N_{present} is current resident count.

7 – Pronunciation

The word “limiting” is essential: it is not the average capacity.

8 – Units

People minus people gives people.

9 – Convention

The supported value must already include the chosen reserve and contingency rules. Do not add the same margin twice.

10 – Why subtraction

Present residents already consume part of the verified capacity; the difference is the remaining person-equivalent headroom.

11 – Assumptions

Per-person demand remains inside the scenario used to derive supported population.

12 – Unit check

person-person=person.

13 – Numerical case

Current population: $N_{\text{present}} = 80$.

Verified capacities are 103, 96, 91, 108 and 94 people.

Limiting verified capacity = $\min(103, 96, 91, 108, 94) = 91$ people.

$H_{\text{admission}} = 91 - 80$.

$H_{\text{admission}} = 11$ people.

14 – Operations

First choose 69 as the minimum; then subtract the sixty residents already present.

15 – Algebra check

$60+9=69$, the limiting verified ceiling.

16 – Mental estimate

Headroom is less than ten people, obviously much smaller than the 40-person step to the final target.

17 – Interpretation

The settlement cannot defensibly jump from sixty to one hundred under this evidence state. Food logistics is the immediate growth gate.

18 – What it does not prove

Nine-person headroom does not mean nine arrivals should automatically be approved; transport batch size, refuge topology, staffing and $N-1$ capacity may impose a stricter gate.

19 – Sensitivity

If improved storage and inventory control raise food-logistics support from 69 to 82, refuge at 72 becomes the new limiting function. Growth work should then shift to the next bottleneck.

20 – Practice

Guided exercise. Find admission headroom for 80 residents with verified capacities 103, 96, 91, 108 and 94.

Detailed guided correction.

1. Limiting capacity = $\min(103, 96, 91, 108, 94) = 91$ people.
2. $H_{\text{admission}} = 91 - 80 = 11$ people.
3. The 11-person headroom is nominal; growth still requires N-1 margins, staffing, consumables, rescue and public-health readiness.

Autonomous exercise. At 80 residents, nominal limiting capacity is 94 people, but under N-1 conditions the critical utilities support 88, 84, 92 and 86 people. A proposed growth step would add six residents. Apply both nominal headroom and N-1 criteria.

► Autonomous correction – open after attempting the exercise

21 – Mission decision

Approve only the population increment that all mandatory functions can support with verified margin; refuse schedule pressure that treats the hundred-person target as irreversible.

Distribution networks become critical infrastructure

At larger scale, a single central utility room connected by long undifferentiated trunks becomes a common-cause hazard. Power, water, data and ventilation should be sectionalised so a fault can be isolated while essential service is rerouted. The network review should identify valves, breakers, cross-ties, maximum transfer capacity and the loads that cannot be served after each credible isolation.

Primary-source bridge. NASA's architecture components include Infrastructure Support, Communications/PNT, Data Systems and Logistics as distinct but interacting sub-architectures. [NASA — Moon to Mars Architecture Components](#).

Sectionalisation has a human cost: more components, inspections, procedures and possible misconfiguration. The architecture therefore needs a balance between isolation capability and operational complexity. A hundred-resident settlement that depends on dozens of manually remembered valve states is not mature infrastructure.

Public health and occupational health become population systems

With one hundred residents, a cluster of gastrointestinal illness, respiratory symptoms or chemical exposure can affect work capacity and critical staffing. Health capability must therefore include surveillance, sanitation, food and water traceability, isolation space, exposure records and outbreak investigation. These functions appear before a large “hospital” because prevention and detection protect the entire workforce.

Primary-source bridge. NASA's Human Research Program develops methods to protect astronaut health and performance for exploration missions. [NASA — Human Research Program](#).

Specialisation creates human common causes

More residents allow dedicated experts, but the settlement can become dependent on a small number of people who hold rare qualifications. Build a skill matrix for every critical function: minimum simultaneous staffing, total qualified personnel, backup discipline, training time, recency and the consequence of losing one or two people. A hundred-person settlement should be able to absorb illness, EVA injury or reassignment without discovering that one technician was the only person who understood a critical process.

Primary-source bridge. NASA HRP Human Factors and Behavioral Performance work provides context for performance, workload and team risks in exploration. [NASA — Human Research Program](#).

Industrial growth should be selected by autonomy leverage

Local industry should not imitate Earth's entire industrial base. Prioritise processes that reduce high-risk imports, shorten recovery time or unlock several other capabilities. Water and gas processing, selected structural parts, seals, simple machining, electrical repair, cleaning agents or feedstock preparation may provide more autonomy leverage than complex products with low failure frequency. Each proposed plant should be reviewed against imported mass avoided, power/heat demand, hazardous inventory, labour, maintenance burden and quality assurance.

Governance must scale before conflict does

A thirty-person outpost can rely heavily on direct conversation. At one hundred residents, work groups, shifts and districts create information boundaries. The settlement needs explicit procedures for safety rules, resource restrictions, maintenance priorities, dispute resolution and emergency authority. Local teams should have enough autonomy to operate efficiently while respecting settlement-wide constraints such as pressure integrity, fire safety, potable-water status and protected power loads.

Growth review: a 30 → 60 → 100 evidence programme

Plateau 30. Demonstrate stable life support, rescue coverage, core maintenance and the first industrial capabilities. Record actual per-person demand and maintenance labour rather than relying only on design assumptions.

Predeployment for 60. Add capacity before population. Commission new trunks, storage, refuge volume, sanitation, medical stocks and workshop capacity. Test rerouting and one major utility loss with only the thirty-person team present.

Plateau 60. Operate long enough to reveal shift coverage, logistics traffic, inventory turnover and equipment wear. Use the measured data to revise the hundred-person architecture.

Predeployment for 100. Expand network topology, public-health functions, specialist training, spare inventories and local production. Re-run N-1 and evacuation/refuge scenarios with the larger geometry.

Plateau 100. Admission is the end of commissioning, not the end of engineering. Growth should stop if verified margins, workforce or maintenance backlog deteriorate.

Graduation scenario: reject growth when the average looks good

Your settlement has sixty residents. Average utility capacity across six functions appears to support eighty people, and a transport window can deliver twenty more residents now. But verified function capacities are 88 for oxygen, 82 for water, 71 for refuge, 76 for sanitation, 91 for power and 69 for food logistics. A public-relations schedule favours accepting all twenty. The correct analysis refuses to average these numbers: the limiting verified function supports 69, so current headroom is only nine before any extra policy margin. Write the HOLD decision, list the work that could raise the limiting capacities and state which evidence would reopen the admission gate.

Final hundred-person evidence package

The graduation dossier should include: population staging; verified capacity table; N-1 utility analysis; sectional network maps; water/air/food inventories; thermal and power envelopes; public-health plan; medical capability; maintenance backlog and spare strategy; critical-skill matrix; mobility/rescue coverage; industrial road map; governance authority map; Earth-dependence register; construction plan; six-event integrated contingency timeline; and explicit growth gates. A reviewer should be able to trace every “we can support 100 people” statement back to demonstrated capacity, a calculation or a clearly labeled assumption.

Source context. NASA Moon-to-Mars architecture, ECLSS and NASA-STD-3001 provide primary context for human exploration systems. The hundred-person settlement and numerical capacities here are a Delta-Sierra educational architecture, not an official NASA plan. [NASA — Moon to Mars Architecture Components](#).

Utility districts: design isolation and rerouting together

At one hundred residents, power, water and ventilation trunks should be mapped as infrastructure networks. Every isolation point creates both protection and a possible capacity bottleneck. Review the largest district that can be lost without forcing evacuation of the entire settlement. Then verify the remaining network can route protected demand without exceeding feeder, pipe, valve or fan limits.

Cross-ties should not exist only on drawings. They require periodic exercise, compatible pressure or voltage conditions, operator procedures and maintenance. A cross-tie that has never been operated under realistic load is an assumption, not demonstrated resilience.

Sanitation and waste become municipal-scale functions

More residents increase wastewater flow, solid waste, cleaning demand and the number of points where contamination can spread. Sanitation capacity should include peak flow, storage during processor outage, isolation of suspect streams and cleaning supplies. Waste-processing equipment may share heat, power or ventilation with other industrial functions; those dependencies belong in the growth gate.

Internal logistics becomes a traffic system

A larger settlement moves food, parts, samples, waste, medical supplies and construction materials every day. Corridors, airlocks, lifts or carts can become bottlenecks. Separate clean and dirty flows where practical, protect emergency egress and schedule high-volume cargo movement so it does not block routine operations. Inventory location matters: a spare part on the opposite side of a depressurised district may be physically present yet operationally unavailable.

Maintenance backlog as a growth inhibitor

Population growth should stop when preventive maintenance is persistently deferred. Track overdue critical tasks, repeat failures, cannibalisation and work orders waiting for specialist labour. A rising backlog indicates that installed capacity is outrunning the organisation's ability to sustain it. The hundred-person target should not be achieved by borrowing reliability from the future.

Primary-source bridge. NASA's Reliability and Maintainability standard provides primary context for maintainability and reliability discipline. [NASA — Reliability and Maintainability Standard](#).

Training pipeline as infrastructure

At one hundred residents, qualification cannot rely on informal apprenticeship alone. Define curricula, supervised tasks, recency checks and requalification for critical roles. Training consumes expert time, so the growth plan should schedule it before staffing becomes thin. New

arrivals should not count as full operational capacity the day they step out of the lander.

Primary-source bridge. NASA's Moon to Mars architecture includes Human Systems and Infrastructure Support, reinforcing that people and support systems must be planned together. [NASA — Moon to Mars Architecture Components.](#)

Construction while inhabited: manage interference

Expansion creates dust, noise, power draw, temporary isolations, vehicle traffic and pressure-boundary work near people who are already living on site. The growth plan should separate construction zones, schedule utility tie-ins, provide temporary services and protect evacuation routes. A construction project that repeatedly places the operating settlement into degraded mode can erase the safety margin it is meant to create.

Primary-source bridge. NASA's Systems Engineering Handbook provides context for configuration, interfaces, verification and lifecycle integration. [NASA — Systems Engineering Handbook.](#)

Use recovery time as the operational autonomy metric

Apply the progressive-autonomy framework by asking how long the settlement would need to recover a lost capability if Earth stopped supplying the relevant item today. The number combines stock, repair, local manufacture, substitution and transport delay in one operationally meaningful clock.

Primary-source bridge. NASA's Moon to Mars architecture identifies Logistics Systems as a sub-architecture covering packaging, transport, staging, storage, tracking and transfer of cargo. [NASA — Moon to Mars Architecture Components.](#)

Economy of scarce attention

One hundred residents create more capability but also more coordination. Meetings, permits, planning, documentation and conflict resolution consume human time. Governance and software should reduce unnecessary coordination while preserving safety-critical interfaces. Measure administrative load and investigate processes that repeatedly require the same information to be entered or approved.

Final defence: demonstrate a reversible growth decision

The strongest capstone is not one that always says “grow.” Present one stage where the settlement deliberately pauses expansion because a limiting function, skill pool or maintenance backlog fails the admission criterion. Then show the work package that restores margin and the evidence that allows growth to resume. Reversibility is a mark of mature engineering because it prevents schedule pressure from converting an assumption into an irreversible population commitment.

R59 growth-control board: a hundred-person settlement must be able to stop growing safely

Growth is easiest to approve when the average numbers look comfortable: total power exceeds average load, tanks are not empty and beds exist. The harder question is whether the settlement can absorb the next population increment without turning every maintenance event into an emergency. A mature growth programme therefore treats admission as a reversible infrastructure decision governed by verified capacities, N–1 states, staffing, backlog and recovery time.

Use plateaus as operational experiments

Do not move continuously from 60 to 100 residents. Hold at planned plateaus long enough to observe seasonal demand, equipment faults, medical events, maintenance accumulation, training load and social/operational friction. A plateau ends when evidence is mature, not merely when the next transport arrives. Cargo can arrive before people; people should not be used as the commissioning load for unproven infrastructure.

District utilities need isolation without orphaning residents

At larger scale, distribution becomes a network problem. Water, power, data, ventilation and waste transfer may cross district boundaries. Isolation protects the settlement from propagation, but badly chosen isolation can strand a district without minimum services. Map each valve, switch and network cut with the population it can separate, the alternate path available and the time until that district reaches a protected limit.

Maintenance backlog is a population variable

A settlement can remain technically functional while preventive maintenance is deferred faster than teams can recover it. Track backlog in labour-hours and criticality, not only task count. Growth should pause if new residents consume maintenance capacity faster than they add qualified capability. A rising backlog in low-visibility components—filters, seals, connectors, fire barriers, sensors—can be an early sign that headline production capacity is masking institutional overload.

Training pipeline needs lead time

Specialists cannot be created instantly when a new reactor, chemical process, medical capability or utility district is commissioned. For each critical role, record the number of qualified operators, assessors and instructors, how long qualification takes, and what happens if one person is unavailable. Cross-training reduces single-person dependency, but only if practical performance is assessed rather than inferred from reading a procedure.

Public health changes the consequence scale

At one hundred people, a contamination event, respiratory outbreak or sanitation failure can affect enough residents to reduce staffing in multiple technical domains at once. Public-health surveillance, isolation capacity, hygiene resources, waste handling and occupational exposure control should therefore appear in the same growth board as power and water. Health capacity is not a separate annex when workforce availability is part of system resilience.

Track the recovery-time trend before opening the next growth gate

Use the recovery-time metric as a growth trend rather than restating its definition. If population rises while the recovery time of a critical imported capability also rises, the settlement is becoming more fragile even when nominal inventory is larger.

Board exercise — nominally ready, N-1 not ready

Eighty residents are present. Nominal limiting capacity supports 94, so a six-person arrival appears safe. Yet one life-critical N-1 state supports only 84. The growth board should reject the argument that “there is room for fourteen.” After the six-person arrival, population would be

86, exceeding the N-1 support by two. The correct action is HOLD until the limiting utility is upgraded, a valid contingency raises its N-1 capacity, or the arriving population is reduced. The discipline is simple: growth consumes future recovery margin as well as nominal capacity.

Primary-source bridge. NASA's Moon to Mars architecture describes logistics, habitation, human systems, infrastructure support, communications and other sub-architectures as parts of one exploration system. The staged population gates here are a Delta-Sierra teaching architecture, not a NASA plan for a hundred-person Mars settlement. [NASA — Moon to Mars Architecture Components](#).

R60 capstone expansion: operating a hundred-person settlement as a network of districts, professions and recovery paths

At one hundred residents, the settlement changes category. The main engineering problem is no longer merely scaling every tank and generator by population. Distribution networks, district isolation, public health, maintenance governance, training pipelines, internal logistics and institutional workload become first-order design variables. The capstone should therefore defend not only “can the settlement support one hundred people?” but “can it remain governable, maintainable and recoverable when multiple districts and professional teams interact?”

Design districts around failure containment as well as convenience

A district boundary should help stop faults from propagating. Power feeders, water loops, data networks, ventilation and waste transfer need isolation points, alternate paths where justified and clear consequences when a segment is cut. The map should show which residents and functions lose service after each isolation. A district that can be isolated only by abandoning its medical room or trapping a population without sanitation is not a useful resilience boundary.

Primary-source bridge. NASA's Moon to Mars Architecture describes multiple interacting sub-architectures including habitation, infrastructure support, logistics, communications/PNT and human systems. The district model here is a Delta-Sierra teaching architecture for exploring those interactions at larger settlement scale. [NASA — Moon to Mars Architecture Components](#).

Staffing must be analysed by qualification depth

Headcount is not capability. For each critical profession, record qualified operators, maintainers, independent verifiers and instructors. A system may have six people who can run a normal checklist but only one who can diagnose an abnormal condition or certify another operator. That is a hidden staffing bottleneck. Growth should include the lead time required to produce the next qualified maintainer or instructor, not assume that new residents instantly add equivalent technical depth.

Public health can disable technical systems indirectly

An infectious outbreak, foodborne event or sanitation failure can remove multiple workers from unrelated technical teams at once. Public-health surveillance therefore belongs on the same resilience board as power and water. Track isolation capacity, hygiene supplies, diagnostic capability, workforce absenteeism thresholds and the technical functions that lose coverage if a cluster affects one shift or profession. The consequence is systemic even when the health event begins locally.

Primary-source bridge. NASA's Human Research Program studies health and performance risks for exploration. The population-health board in this capstone extrapolates that concern to a much larger resident community. [NASA — Human Research Program](#).

Internal logistics becomes a transport network

At thirty people, a technician may personally know where most critical spares are stored. At one hundred, the settlement needs addresses, custody, controlled environmental storage, retrieval priorities and movement rules. Time to move a filter, pressure vessel, medical kit or repair team between districts becomes part of recovery. The capstone should therefore measure internal delivery time under normal and degraded routes, not only whether the item exists somewhere in inventory.

Maintenance governance should expose backlog by risk

Backlog needs more than a task count. Separate overdue preventive work, corrective work, inspection findings, temporary repairs and deferred upgrades. Rank them by consequence, remaining redundancy and required skill. A thousand cosmetic tasks can be less dangerous than one overdue inspection on a common-cause isolation valve. Growth should pause when the backlog is consuming protective depth faster than teams can restore it.

Primary-source bridge. NASA's Reliability and Maintainability standard provides context for designing and managing systems with reliability and maintainability objectives. [NASA — Reliability and Maintainability Standard](#).

Utility operators need cross-district command rules

A local district operator may be able to protect one area by actions that harm another. Examples include shedding power, closing a water cross-connect, isolating data, reducing ventilation or reserving a rover. The governance model should define when local authority is sufficient and when the decision becomes settlement-wide. The escalation rule should be tied to consequence and time-to-limit so that cross-district coordination does not create paralysis.

Growth gates should test the next credible failure, not only nominal capacity

Before admitting another population increment, rerun the capacity ledger under a named degraded state. If nominal water supports 118 residents but the credible N-1 water configuration supports 96, admitting a population of 100 consumes the failure margin even though the nominal number looks comfortable. A defensible gate asks whether the settlement can sustain the new population through the failure states that the architecture claims to tolerate.

Prioritise imported dependencies by recovery delay and consequence

Rank imported dependencies on two axes: time to restore the capability after loss and consequence while waiting. The resulting list tells the industrial programme where local repair, alternate suppliers, substitution or stock depth buys the largest autonomy gain.

Economic activity must not hide life-support cross-subsidies

A workshop can appear productive because it receives power, cooling, water treatment, logistics and maintenance from settlement-wide systems that are not charged to its local budget. For architecture decisions, expose those dependencies physically even if no monetary accounting system is used. A process that produces valuable output but consumes the only protected maintenance capacity or thermal margin may be strategically harmful during that period.

Hundred-person integrated exercise — three simultaneous pressures

Run a capstone in which one district loses a utility feeder, the medical team reports a cluster of gastrointestinal illness and the maintenance backlog contains an overdue isolation-valve inspection. The team must decide whether to isolate the district, move residents, reserve sanitation resources, reassign staff and postpone industrial work. The evaluation should track protected services, staffing, cross-district consequences, communication and recovery debt. The point is to force engineering, public health and governance into one decision space.

Final growth board — evidence for 100 → 110 should be harder than 80 → 90

As population rises, the settlement should demand more mature evidence because consequences and coupling increase. A ten-person increment at one hundred people can require new shift coverage, sanitation capacity, instructor depth, medical isolation, district logistics and utility redundancy even if average per-person consumption remains unchanged. Growth is therefore not linear scaling. The board should be willing to HOLD at 100 until organisational and recovery capacity catches up with hardware capacity.

Primary-source bridge. NASA's Moon to Mars Strategy and Objectives describes a structured approach to exploration goals and architecture. The staged 100-person growth gates here are Delta-Sierra educational constructs, not an official NASA Mars settlement plan. [NASA — Moon to Mars Strategy and Objectives](#).

R60 hundred-person defence: prove that growth does not outrun institutions

Shift coverage and professional depth

Create a seven-day duty roster for power, life support, medical response, communications, EVA/rescue and industrial safety. Then remove one specialist from each of two professions. The settlement should still have a documented response for night shifts, maintenance periods and concurrent incidents. A population of one hundred can create an illusion of abundant labour while critical qualifications remain concentrated in a handful of people.

District isolation and resident movement

Choose one district and simulate loss of water or power for twelve hours. Determine how many residents must move, where they can be received, what sanitation and sleeping capacity exists, how medicines and personal equipment follow them, and whether the receiving district remains within its own protected capacity. Emergency relocation is itself a load case.

Public-health staffing shock

Assume an illness removes fifteen percent of residents from normal duty for three days, with cases concentrated in one work group. The capstone should identify which technical functions lose coverage, which work is deferred and which cross-trained personnel can substitute. This tests the coupling between health and infrastructure more realistically than adding a generic “medical margin.”

Logistics congestion

Simulate simultaneous demand for a rover, lifting equipment and two technicians by three different recovery tasks. The team must prioritise shared assets and show the consequence of delaying each job. At larger scale, logistical conflict becomes a system property even when enough total equipment exists.

Governance under growth pressure

Introduce a scheduled arrival that has political or mission pressure to proceed even though one growth gate is marginal. The defence should state the evidence threshold for HOLD and resist substituting schedule commitment for technical readiness. The ability to stop growth is one of the clearest demonstrations that the architecture treats population as a safety-critical load rather than a prestige metric.

Primary-source bridge. NASA’s Systems Engineering Handbook emphasises verification, validation, risk and lifecycle thinking. The institutional stress tests above extend those disciplines into the operational design of a large Mars settlement. [NASA — Systems Engineering Handbook](#).

R60 metropolitan-scale preparation: the first hundred residents are

the point where hidden institutions become infrastructure

At this scale, competence databases, maintenance planning, inventory control, public-health surveillance, incident learning, training records and configuration management are not administrative extras. They are infrastructure because technical systems depend on them to remain safe. The capstone should therefore assign ownership, staffing and failure modes to these information processes just as it does to pumps and power buses.

Consider qualification records. If the database is wrong, a shift can appear covered when the only current operator is actually expired or unavailable. Consider inventory: a duplicated record can create a false spare. Consider configuration management: a procedure written for an old valve layout can direct a crew to isolate the wrong branch. These are information failures with physical consequences. Redundancy and backup of the underlying data therefore belong in the resilience architecture.

Training throughput limits growth

Growth can require more qualified people faster than the settlement can train them. For each profession, estimate instructor capacity, supervised practice slots, assessment time and the minimum experience required before independent duty. A surge of new residents can temporarily reduce net productivity because experienced staff spend time teaching. The growth board should model that transition rather than assuming every arrival immediately adds full labour capacity.

Governance needs a hierarchy of decisions

Routine local decisions should remain local; settlement-wide decisions should escalate. Define categories such as local reversible action, cross-district resource reallocation, emergency reserve consumption, population movement and strategic growth. Each category needs an authority level and a fallback if the nominal authority is unavailable. Without this hierarchy, either everything escalates and the organisation becomes slow, or too much remains local and one district can consume shared protection without coordination.

Incident learning must protect against recurrence without paralysing operations

After a significant anomaly, record the technical cause, contributing organisational conditions, detection path, recovery actions and what evidence proves the corrective action works. Not every event requires a major redesign, but repeated “operator error” labels should trigger deeper investigation into interfaces, procedure clarity, workload, training or automation. A growing settlement should become more knowledgeable after failures, not merely more experienced at surviving them.

Final evidence exercise — show one capability that improves with scale and one that degrades

The capstone team should identify a function that genuinely benefits from one hundred residents—perhaps professional redundancy, medical coverage or maintenance specialisation—and another that becomes harder, such as coordination, infection control, logistics congestion or configuration discipline. Then explain what architecture converts the second effect from an uncontrolled cost of scale into a managed system. This forces the student to abandon the assumption that “more people” is uniformly beneficial or uniformly burdensome.

Training throughput, professional depth, governance and maintenance backlog become infrastructure that can limit safe population growth. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Primary sources and bridges

- [NASA — Moon to Mars Architecture Components](#)
- [NASA — Environmental Control and Life Support Systems](#)
- [NASA-STD-3001 Volume 2](#)