

Settlement power: budgets, microgrids and priorities



Conceptual visualization — a microgrid is an operational network of generation, storage, distribution, protection and prioritized loads.

Size a settlement where generation, storage, distribution and load shedding must continue through dust, maintenance and source loss.

Mastery objectives

- build separate instantaneous-power and time-integrated-energy budgets for a settlement microgrid
- classify loads by consequence, interruption tolerance and restart dependency before shedding them
- evaluate storage, dispatchable generation, protection and black-start paths under degraded conditions
- prove that restoration leaves adequate reserve, stable power quality and an auditable final configuration

1. Peak power and daily energy answer different questions


A system can have enough energy over a day while being unable to supply a short peak. Budgets therefore track instantaneous kilowatts and accumulated kilowatt-hours. Pumps, furnaces, communications and vehicle charging cannot simply be added as if all loads operated with the same time profile.

2. Classify loads before calculating generation

Vital, critical, deferrable and discretionary loads do not have the same claim on power. Life support, minimum thermal control and safety communications remain high priority. Manufacturing, electrolysis or charging may sometimes be shifted. The hierarchy is written before the emergency so that shedding is not improvised under pressure.

3. Microgrid: multiple sources, multiple stores, one stability problem

A base may combine solar, fission, batteries and thermal storage. The microgrid shares power while maintaining voltage, frequency or DC-bus limits. Loss of one converter must not turn a local failure into a base-wide blackout.



A microgrid is a coupled network; reserve and protection depend on how sources, storage and loads share buses. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

4. Storage: reason in usable autonomy, not nameplate capacity

A 1 MWh battery does not necessarily provide 1 MWh to loads. Depth of discharge, temperature, aging, efficiency and safety reserve reduce usable energy. The budget states the energy that can actually be delivered in the critical scenario.

5. Dust, night and maintenance: nominal generation is not available 100% of the time

Solar production and nuclear availability both change with environment and maintenance. Energy architecture therefore uses time profiles and failure scenarios. Storage bridges mismatches between demand and generation; it does not excuse chronic undersizing.

6. Distribution and common-cause failures

Two sources feeding the same converter or running through the same cable are not truly independent. Redundancy separates common causes across location, conversion, protection, cooling and control.

7. Microgrid stability: balance power and energy

A positive 24-hour energy balance does not guarantee a stable grid. A base can produce enough total energy and still lack instantaneous power during compressor startup or immediately after loss of a source. Energy in kilowatt-hours must be distinguished from power in kilowatts. Storage has to absorb transients and give dispatchable sources time to take over. Control strategy also needs the priority and dynamics of every load: a greenhouse may be reduced briefly, while a medical pump may require uninterrupted supply.

8. Black start: rebuild the grid from almost nothing

After a total outage, many components required for restart themselves need electricity: controllers, pumps, heaters, communications and converters. Black-start capability means recreating supply without relying on an already energized grid. An autonomous source establishes the first bus, then branches reconnect in an order that avoids simultaneous inrush loads. The procedure should be tested with cold batteries, converters that are actually available, and explicit criteria for authorizing each reconnection step.

9. Energy growth: do not consume the entire margin

A growing settlement adds greenhouses, workshops, laboratories, communications and living areas. If each expansion consumes all available reserve, failure margin disappears. The budget should distinguish nominal demand, peak demand, operating reserve, contingency reserve and storage energy that remains usable after ageing and depth-of-discharge limits. Electricity must also be linked to heat: a machine rejecting several kilowatts can be useful in one condition and a cooling burden in another. Flexible loads should be scheduled rather than sizing every source for the worst simultaneous case.

10. Worked example: battery autonomy in refuge mode

Refuge vital loads consume 18 kW. A battery has 420 kWh nameplate capacity, but only 80% depth of discharge is allowed and end-to-end efficiency is 92%. Usable energy = $420 \times 0.80 \times 0.92 = 309.1$ kWh. Autonomy = $309.1/18 \approx 17.2$ h. A 24 h requirement is therefore not met even though the battery is advertised as 420 kWh.

Calculated case study: shedding a microgrid without losing life-critical loads

TEACHING ASSUMPTION — Colony loads are 22 kW critical, 18 kW important and 35 kW deferrable. A fault reduces available generation to 38 kW.

Let P_c be critical power, P_i important power, P_d deferrable power and P_g available generation, all in kW. A kilowatt measures instantaneous power; energy over time would be expressed in kWh.

Total pre-fault load: $P_{\text{tot}} = 22 + 18 + 35 = 75$ kW. After supplying the 22 kW critical load, $38 - 22 = 16$ kW remains. The important category is short by $18 - 16 = 2$ kW, and all 35 kW of deferrable load must be shed. Total shedding is $35 + 2 = 37$ kW, or $37 \div 75 \approx 49.3\%$ of the initial load.

The calculation shows why a priority list must exist before a fault. Having 38 kW is not enough: operators must already know which loads can stop, for how long, and what restart consequences follow.

11. Exercise

Build a 24 h budget with 25 kW vital loads, 40 kW critical loads and three deferrable loads. Simulate loss of 50% of generation for 18 h and create a load-shedding order.

12. Reasoned solution

First classify loads as immediate survival, operational continuity or deferrable, then compare survival power with the remaining source. If it exceeds available capacity, priorities or independent supply must change; shedding only comfort loads cannot solve a structural deficit.

13. Validation mini-project

Design a Mars microgrid: sources, buses, protection, storage, load profiles, priority classes, failure scenarios, restart strategy and stability criteria.

Associated calculators

Apply this module with:

- [Mars solar + battery calculator](#)

A settlement power system must survive the instant, the day, the dust event and the repair window

A power budget answers several different questions that must not be collapsed into one number. Peak power asks whether generation, converters and distribution can support the highest simultaneous demand. Daily energy asks whether enough kilowatt-hours can be produced and stored over time. Reserve asks what margin remains after uncertainty and failures. A design can pass the daily energy calculation and still collapse during a short motor start, or it can meet the peak and slowly drain its batteries because average production is too low.

Mars adds operational variability. Solar output changes with time of day, season, dust loading, panel orientation and atmospheric conditions. Storage loses usable capacity with temperature, aging and operational limits. Dispatchable sources may have maintenance outages or fuel constraints. The microgrid therefore needs a dispatch philosophy: which source carries steady load, which follows variation, what state-of-charge reserve must be preserved, and which loads can be delayed or shed without endangering the crew.

Load categorization is the foundation of safe shedding. Life support, thermal protection, essential communications, medical capability and control systems belong in a different class from discretionary manufacturing or comfort loads. The list should be more precise than “critical/noncritical.” Some loads can be interrupted for minutes but not hours; others must finish a process safely before shutdown. Priority logic should reflect consequence, restart difficulty and the state of other systems.

Black start deserves its own architecture. After a total or near-total collapse, the grid cannot assume that every controller, pump and network service is already powered. A small protected source may need to energize control, communications, thermal management and one generation path in a deliberate sequence. Each step consumes power before the full system is available. A settlement that can operate normally but cannot restart from darkness has an incomplete energy design.

Growth turns energy margin into a governance question. New laboratories, greenhouses, machine tools and habitats compete for limited reserve. If every expansion consumes the current spare capacity, the colony becomes fragile even while nominal production increases. Power planning should therefore preserve contingency margin and identify the next capacity upgrade before the current system reaches its operational ceiling.

Twelve ideas for reading a microgrid like an operator rather than a brochure

1. Classify loads by consequence and allowable interruption

Every load should have a reason for its priority. A carbon-dioxide removal fan may be immediately important, a freezer may tolerate a short outage but not many hours, and a machine tool can usually stop after reaching a safe condition. The classification records maximum interruption time, restart behavior, surge demand and whether the load has a local backup. This creates a usable load-shedding table instead of a vague list of “critical equipment.”

Priorities can change with mission state. A medical imaging device becomes critical during surgery; a laboratory freezer becomes more important if it contains irreplaceable samples; a greenhouse heater may become time-critical during a cold excursion. Operational software

should therefore allow controlled priority changes with authorization and logging rather than hard-coding one permanent order that ignores context.

2. Apply the power/energy distinction to the load schedule

Power is the rate at which electrical work is delivered at an instant, while energy accumulates power over time. A 20 kW load running for one hour uses 20 kWh; the same load for ten minutes still requires 20 kW capability but only about 3.3 kWh. Cables, converters and generators care about the peak; storage autonomy and fuel or solar collection care about energy.

Operators need both views on the same timeline. A daily total can hide coincident peaks created by cooking, charging rovers and running industrial equipment at the same time. Scheduling flexible loads can reduce the peak without reducing total useful work. That is often cheaper and more resilient than adding hardware solely to support avoidable simultaneity.

3. Build a load profile rather than multiplying one average number

A load profile places expected demand against time: base life-support load, scheduled work, thermal cycles, charging periods and uncertain peaks. The profile reveals how much reserve exists at each hour and where storage must absorb mismatch between production and consumption. It also allows the team to test abnormal days, such as a medical event coinciding with a rover recharge.

Profiles should be updated from measured data after deployment. Actual duty cycles may differ from design assumptions because crews change routines or equipment ages. Comparing prediction with telemetry makes the budget a living operational tool. A persistent difference is not “noise”; it is evidence that planning assumptions or equipment condition need review.

4. Solar production is an envelope, not a fixed nameplate value

Panel nameplate power is measured under defined conditions that do not represent every Martian hour. Sun angle, atmospheric dust, deposited dust, temperature, wiring losses and power electronics reduce delivered energy. A robust plan uses conservative production envelopes and tracks actual array performance over time. Cleaning or redeployment becomes an operational task with a measurable energy benefit.

Solar uncertainty also couples to storage. A low-production day may be harmless if state of charge is high, but the same day after several poor days can cross a reserve threshold. Forecasting therefore needs both expected incoming energy and current stored energy. Operators manage a trajectory, not a single percentage printed next to a solar icon.

5. Dispatchable generation provides control but still needs availability accounting

A dispatchable source can be commanded when needed, which makes it valuable for stabilizing a variable grid. But it is not automatically available 100% of the time. Maintenance, thermal limits, fuel supply, startup time and component failures create outages. The energy plan should distinguish installed capacity from capacity that is credibly available during the scenario being tested.

If one unit is the only path to black start or to winter survival, its maintenance schedule becomes a colony-level risk. Diversity, spare parts and the ability to share load among smaller units can improve resilience. The correct architecture depends on mass, reliability and mission context, but the single-point dependency must be visible.

6. State of charge is an operating boundary, not just a battery gauge

Batteries are usually not planned around using every theoretical kilowatt-hour. Upper and lower state-of-charge limits protect life, power capability and contingency reserve. The usable window may shrink with temperature or aging. Operators should therefore distinguish nominal capacity, currently available capacity and the portion intentionally reserved for emergencies.

A state-of-charge target can also guide dispatch. If forecast production is uncertain, the system may keep a larger reserve before night. If dispatchable generation is healthy, it may accept a lower battery target and run that source later. These choices should be explicit and reproducible rather than improvised by whoever is on shift.

7. Conversion and distribution losses belong in the budget

Energy generated at a source is not identical to energy delivered to loads. Converters, inverters, transformers, cables and battery charge/discharge cycles introduce losses. Some losses appear as heat that the thermal system must remove, linking electrical efficiency to habitat cooling. A budget should state the boundary at which each power number is measured.

Distribution topology also affects failure propagation. One overloaded bus, shared converter or common cable route can remove multiple sources at once. Physical separation and selective protection may be as important as total installed generation. Resilience depends on where power can still flow after a fault, not merely on adding all generator nameplates together.

8. Reserve margin buys time for uncertainty and recovery

Reserve is capacity intentionally left unused during normal planning so the grid can absorb forecast error, a failed source or an unexpected load. The correct margin depends on consequence and recovery time. A system with rapid repair and several independent sources can accept a different margin from a settlement that depends on one difficult-to-service generator.

The margin should be visible in operational dashboards and expansion decisions. If routine operations continuously consume the reserve, the settlement is effectively operating in a degraded state even if no alarm is active. Restoring margin should then become a planned task, just like repairing a known hardware defect.

9. Load shedding should be staged, predictable and reversible

Shedding works best when operators know exactly what will disappear at each threshold. The first stage may pause vehicle charging and industrial heating, a later stage may reduce comfort conditioning, and only severe states should threaten functions close to life support. Each stage should estimate the power recovered and any secondary consequences, such as a process that needs a controlled shutdown.

Automatic shedding can be faster than a human response, but it must remain understandable. A crew facing darkness should not discover that an algorithm disabled a pump for reasons no one can reconstruct. Clear thresholds, event logs and manual recovery authority make automation a tool rather than a new source of uncertainty.

10. Black start is a sequence of dependencies

After a broad outage, the grid may need power to start the very equipment that creates power. Control electronics, cooling pumps, valves, communications and excitation systems can all require an initial source. A protected battery or small independent generator can provide this seed energy, but its capacity must cover the entire startup sequence with margin.

The procedure should specify which bus is energized first, what health checks occur, when additional generation is synchronized, and which loads remain locked out until stability is demonstrated. Black-start drills can reveal hidden dependencies such as a network switch or sensor supply that normal operation never exposes because it is always already on.

11. Islanding can limit the size of a fault

A microgrid can be divided into electrical islands that operate independently when a common bus is damaged or unstable. Islanding sacrifices some flexibility but can preserve a habitat refuge, medical area or control center while the larger network is repaired. The design needs appropriate switching, local regulation and enough generation or storage inside each protected island.

Operators must know the limits of each island. A bus that is stable at normal load may collapse if someone reconnects a large charger before sources are synchronized. Reconnection therefore follows a controlled sequence with voltage, frequency or other relevant checks. Electrical compartmentalization is analogous to closing a hatch during depressurization: it contains propagation.

12. Growth planning protects future reliability

New users are not simply added to the old peak. Expansion should include load diversity, new coincidence patterns, cable and converter capacity, storage reserve, maintenance burden and black-start implications. A greenhouse may add a smooth base load while a workshop adds highly variable peaks; equal annual energy does not mean equal grid impact.

The planning rule should identify a trigger for the next upgrade before reserve is exhausted. This avoids the pattern of waiting until normal operations routinely approach limits and then treating every new project as an emergency. Energy infrastructure should lead colony growth enough to preserve recovery margin.

Microgrid calculation laboratory: load energy and storage autonomy

Formula 1 — daily energy from a time-based load schedule

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

Starting question

How much electrical energy do scheduled loads consume over one day?

Read aloud

Read: "daily energy equals the sum of each load power multiplied by its operating time."

Symbols, pronunciation and meaning

P_i is power of load i ; t_i is its operating duration; Σ means add the contributions; E_{day} is total energy.

Units

Using kilowatts and hours gives kilowatt-hours.

Origin and status of values

Power comes from measured or specified operating values; duration comes from the planned duty cycle. Starting surges are checked separately because this energy equation averages them over time.

Why this operation

Multiplication converts a rate of energy use into an amount for each operating interval. Summation combines independent load contributions.

Substitution and calculation

A 12 kW base load runs 24 h, an 8 kW workshop runs 5 h, and a 6 kW charger runs 4 h. Energy = $12 \times 24 + 8 \times 5 + 6 \times 4 = 288 + 40 + 24 = 352$ kWh/day.

Calculator entry

Enter $12 \times 24 + 8 \times 5 + 6 \times 4$. Keep each term visible so the schedule can be audited.

Mental estimate

The base load alone is about 300 kWh, so a total a little above 350 kWh is plausible.

Independent check

Average power implied by 352 kWh/day is $352/24 \approx 14.7$ kW. That is higher than the 12 kW base load because the flexible loads add energy during part of the day.

Physical or operational interpretation

The grid must collect at least this energy plus conversion/storage losses over the planning period, while still meeting the higher instantaneous peak when loads overlap.

Plain-English translation

In plain language: the settlement uses about 352 kilowatt-hours during the model day, even though its power at any one moment can be much lower or higher than 14.7 kW.

Variation / sensitivity

Moving the workshop to another hour does not change its 40 kWh contribution, but it can reduce the coincident peak if charging occurs at a different time.

Limit / assumption

This simplified sum assumes constant load power during each interval. Real profiles include variable duty cycles, startup surges, efficiency losses and uncertain operations.

What this does not prove

Daily energy does not establish peak power, startup current, ramp rate, bus capacity, generator adequacy or the timing of coincident loads. Two schedules can have the same kilowatt-hours and radically different instantaneous power requirements.

Boundary case to test

Give one load zero duration and its contribution must vanish. Move a fixed-energy load to another hour and total daily energy should stay unchanged even though the peak may change; this is a useful check that energy and power are not being confused.

Formula 2 — battery autonomy inside an allowed state-of-charge window

$$t = C_{\text{nom}} \times f_{\text{usable}} \times \eta / P_{\text{load}}$$

Starting question

How long can a battery support a specified load without crossing the planned lower state-of-charge limit?

Read aloud

Read: “time equals nominal energy capacity times usable fraction times efficiency, divided by load power.”

Symbols, pronunciation and meaning

C_{nom} is nominal stored energy; f_{usable} is the allowed fraction of that capacity; η is delivered-energy efficiency; P_{load} is electrical load; t is autonomy.

Units

If C is in kWh and P is in kW, time is hours. Fractions and efficiency are dimensionless.

Origin and status of values

Nominal capacity comes from the installed battery; usable fraction reflects state-of-charge limits and current condition; efficiency and load should use demonstrated or conservative operating values.

Why this operation

Multiplying capacity by usable fraction and efficiency estimates energy that can actually reach the load. Dividing by power converts available energy into duration.

Substitution and calculation

For a 500 kWh battery, 60% usable window, 90% delivery efficiency and 30 kW refuge load: delivered usable energy = $500 \times 0.60 \times 0.90 = 270$ kWh; $t = 270/30 = 9$ h.

Calculator entry

Enter $500 \times 0.60 \times 0.90 \div 30$.

Mental estimate

Half of 500 is 250 kWh; after allowing a little more usable fraction and some losses, about 270 kWh is reasonable. At 30 kW that should be about nine hours.

Independent check

Multiply $9 \text{ h} \times 30 \text{ kW} = 270 \text{ kWh}$ and recover the delivered usable energy.

Physical or operational interpretation

Nine hours is not a promise of survival; it is the electrical window for the stated refuge load and battery condition. Other resources can become limiting earlier.

Plain-English translation

The protected battery can carry the 30 kW refuge configuration for roughly nine hours before reaching the planned lower limit.

Variation / sensitivity

If refuge load falls to 20 kW, autonomy becomes 13.5 h. If usable capacity degrades to 50%, autonomy falls to 7.5 h at 30 kW.

Limit / assumption

The equation assumes approximately constant power and efficiency. High currents, temperature, aging and power-

electronics limits can reduce usable capacity or maximum deliverable power.

What this does not prove

The autonomy equation does not prove that the battery can safely deliver the requested power. Current limits, temperature, state of health, inverter limits and cell imbalance can make a nominal energy inventory unavailable at the required rate.

Boundary case to test

If the usable fraction is zero, autonomy must be zero. Reducing load should increase time inversely, but letting P_{load} approach zero makes the simple quotient grow without bound; real self-discharge, housekeeping loads and minimum operating power prevent that mathematical limit from representing indefinite survival.

Grid cases: practice decisions before they become emergencies

A dust event begins with batteries already below target

The first task is to distinguish a forecast problem from an immediate power-balance problem. Current generation, state of charge, expected dust duration and essential load define how much time exists. The crew can stop flexible industrial work, delay vehicle charging and move processes into lower-power states while preserving life support and thermal control. If a dispatchable source exists, starting it early may protect battery reserve rather than waiting until state of charge becomes critical.

The decision should include recovery uncertainty. Dust can last longer than forecast and cleaning can be delayed by EVA constraints. A conservative plan protects enough energy to survive a worse-than-expected period and to perform black start if the grid later collapses. The correct objective is not to minimize fuel use or maximize solar utilization; it is to maintain a credible path through the whole disturbance.

A high-power workshop load trips a shared bus

Operators isolate the faulted branch and confirm whether the trip was caused by genuine overload, a short circuit or protection miscoordination. Reclosing without diagnosis can repeat the event and may damage equipment. The load profile is reviewed to see which other loads were coincident. If the workshop can be rescheduled or soft-started, the same useful work may be possible without increasing installed generation.

The incident also tests distribution architecture. If the shared bus trip removed unrelated vital loads, the event exposes an excessive common-cause dependency. Corrective action may involve selective protection, bus segmentation or moving a vital load to a different feed. A microgrid failure is therefore an opportunity to improve topology, not merely reset a breaker.

One dispatchable generator is unavailable for scheduled maintenance

The maintenance window should be approved only after checking forecast production, battery reserve, remaining generation and credible simultaneous faults. If the unavailable unit is also the preferred black-start source, an alternate startup path must be demonstrated before work begins. The maintenance job itself may consume power for tools, heaters or test equipment, which belongs in the temporary load plan.

A successful maintenance plan states conditions for aborting the work. If dust intensifies, another source develops an anomaly or state of charge falls below a threshold, the team may restore the unit to a safe operable configuration rather than continuing. This makes maintenance part of system operations rather than a separate technical activity.

A medical emergency creates an unexpected priority load

Medical equipment can change both power and duration requirements without warning. Operators should know which discretionary loads can be shed immediately and whether the medical area has local ride-through energy. If surgery or imaging becomes time-critical, priority tables are deliberately updated and logged so automation does not later remove the load because it was classified as normally nonessential.

The energy team also watches secondary systems: thermal removal for medical equipment, lighting, data, sterilization and environmental control. The useful lesson is that load priority attaches to mission consequence, not to equipment identity. The same device can move from low priority to life-critical depending on context.

A total outage requires black start

The crew follows a rehearsed sequence from the protected seed source. Essential controls and communications are energized first, then the startup support needed by one generation path. The team verifies bus stability before adding life-support groups and keeps large flexible loads locked out. Each step checks that the next dependency is available rather than assuming the whole network will recover at once.

If a step fails, the procedure identifies a branch point instead of forcing repeated attempts that drain the seed battery. After stable generation returns, storage is recharged and loads are restored by priority. The post-event review compares actual startup energy and time with the design assumptions, because black-start performance must be measured just like any other emergency capability.

A new greenhouse asks for the last 15% of apparent spare capacity

Approving the load solely because normal telemetry shows 15% spare power would consume the contingency margin. Planners should test peak coincidence, daily energy, seasonal production, storage reserve and the failure of the largest source. If the greenhouse is biologically important, its interruption tolerance and thermal inertia also matter. The expansion may justify new generation or storage before occupancy grows dependent on the crops.

This case turns reserve into a governance rule. A settlement should define how much margin cannot be allocated to routine growth without an explicit risk decision. Otherwise each locally reasonable project gradually removes the system-level ability to absorb faults.

Microgrid dispatch drills with worked reasoning

Exercise 1 — Separate power from energy

A 25 kW machine operates for 2 h while a 10 kW base load runs all day. State the peak if they overlap and calculate total daily energy from those two loads.

► Reveal the reasoned solution

Exercise 2 — Reschedule a flexible load

A 12 kW rover charger and an 18 kW workshop currently run at the same time on top of a 14 kW base load. What peak can be avoided by scheduling them separately?

► Reveal the reasoned solution

Exercise 3 — Calculate refuge battery time

A 600 kWh battery allows 50% of capacity for refuge service, with 90% delivery efficiency. Refuge load is 27 kW. Estimate autonomy.

► Reveal the reasoned solution

Exercise 4 — Design the first shedding stage

Name four loads you might shed before life support during an energy deficit and explain the selection logic.

► Reveal the reasoned solution

Exercise 5 — Test a maintenance outage

The largest generator will be offline for six hours. List the evidence required before approving the job.

► Reveal the reasoned solution

Exercise 6 — Explain islanding

A short circuit damages the central bus but a habitat has local storage and a switchable feeder. What is the purpose of islanding?

► Reveal the reasoned solution

Exercise 7 — Protect reserve during growth

Normal operations use 82% of credible available generation at peak. A new facility requests another 12%. Why is simple arithmetic not enough to approve it?

► Reveal the reasoned solution

Exercise 8 — Build a black-start sequence

Put these actions in a defensible order: restore large workshop, energize protected control bus, start one stable generator, verify distribution, restore essential life support.

► Reveal the reasoned solution

Microgrid incidents: dispatch, shedding and recovery decisions

A battery string develops one weak module

The pack may still show substantial total energy while the weak module limits current or usable depth of discharge. Operators compare cell or module trends, temperature and balancing behavior and reduce the allowed operating window if necessary. Continuing to plan from nominal pack capacity can create a sudden loss of reserve when the weak element reaches its limit first.

Maintenance planning considers whether the string can be isolated, whether spare modules are compatible and how replacing one element affects balancing. Storage health is therefore a configuration-dependent capability, not a fixed kilowatt-hour label.

Solar cleaning requires an EVA during tight power conditions

The team compares the energy expected from cleaning with the energy and crew-risk cost of the EVA. If dust loading is moderate and a dispatchable source can bridge the deficit, delaying the EVA may be safer. If state of charge is falling toward a critical threshold, cleaning may become time-sensitive. The decision is based on predicted energy recovery and survival margin rather than a routine calendar.

Robotic cleaning or panel orientation may reduce future EVA dependence. The operational history should capture how much production was actually restored by cleaning, allowing later planners to estimate benefit instead of relying on assumptions.

A converter runs hot near its rated load

Electrical capacity cannot be separated from thermal rejection. A converter that is electrically within rating may derate when cooling is poor or ambient temperature is high. Operators review temperature trend, cooling flow, dust accumulation and load distribution and may shift flexible

demand to another bus. The thermal system receives the converter losses as heat, so a hot electrical day can also be a cooling day.

Design review asks whether redundant converters share the same cooling loop. If one thermal fault disables both electrical paths, nominal redundancy is weaker than it appears. Common support dependencies belong in the microgrid reliability model.

A load-shedding event repeatedly interrupts the same scientific freezer

The freezer may have been classified as deferrable during initial design, but accumulated irreplaceable samples can change its consequence. Priority tables should therefore support controlled reclassification and local ride-through storage where justified. Operators document why the priority changed and what other load now sheds first.

This prevents a static automation rule from quietly destroying high-value assets. Electrical governance must track mission context, not only equipment type.

Night begins with an unexpected industrial process still running

The controller should not assume the process can be cut instantly. Operators determine whether stopping now damages material, traps hazardous heat or makes restart more energy-intensive than finishing the batch. If safe, the process is paused; if not, other flexible loads may shed while the team calculates the energy needed to reach a safe stop point.

This is why load classification includes shutdown behavior. A “noncritical” machine can temporarily become operationally important because an unsafe interruption creates a secondary hazard.

A distribution feeder fault isolates a habitat wing

Local storage and generation inside the wing can support an island if switching and regulation were designed for it. Operators first ensure the faulted feeder is isolated and cannot be re-energized inadvertently, then establish the island load list and duration. Nonessential loads remain locked out because local capacity is smaller than the normal grid connection.

Reconnection follows explicit checks after repair. Simply closing the feeder when both sides are energized can create damaging transients if conditions are not compatible. Islanding is useful only with disciplined separation and restoration.

A generator start fails twice during a dust event

Repeated starting attempts can consume battery energy, fuel or starter life. The procedure limits retries and directs diagnosis of permissives, thermal state, fuel path and electrical support. Meanwhile the energy team recalculates how long remaining storage can carry critical loads and may move immediately to a deeper shedding stage.

The key decision is whether troubleshooting can finish before reserve reaches its protected threshold. If not, the crew transitions to a survival configuration rather than spending the last energy on uncertain start attempts.

A new machine tool introduces short high-current pulses

Daily energy may be small while pulse power stresses converters and causes voltage dips for sensitive electronics. Measurements capture the waveform and coincidence with other loads. Soft-start hardware, local buffer storage or scheduling can solve the problem without a major increase in total generation.

This case shows why integrating industrial equipment requires power-quality review, not only adding its kilowatt-hours to the annual budget. Peak and dynamic behavior can dominate compatibility.

A thermal-control failure raises habitat electrical demand

Heaters, pumps or fans may run longer when insulation, radiators or environmental conditions degrade. The energy team sees the higher load as a symptom as well as a grid problem. Shedding unrelated loads can preserve stability temporarily, but correcting the thermal fault is what restores sustainable margin.

Cross-system dashboards should therefore link demand anomalies with physical causes. Otherwise operators may normalize the rising electrical baseline and gradually lose reserve without recognizing the failing subsystem behind it.

A settlement plans a major expansion during its best solar season

Capacity approval should use the worst relevant seasonal and fault conditions, not the month when telemetry looks most generous. Planners model lower production periods, storage aging and maintenance outages and identify the upgrade date before the new population or greenhouse becomes dependent on the capacity.

Growth should also include black-start and islanding tests. New loads can change the sequence and seed-energy requirement even if annual energy remains acceptable. Infrastructure expansion is complete only when abnormal recovery still works.

Power-system design dossier: reserve, protection and restoration

Design dossier — define the electrical survival floor

Before optimizing generation, the settlement should define a minimum electrical configuration that keeps people safe: essential air revitalization, thermal protection, command and communications, critical medical capability, minimal water delivery and whatever instrumentation is needed to monitor those functions. This survival floor is a concrete load list with measured power, not an abstract percentage of normal demand.

Once the floor is known, battery reserve and emergency generation can be sized around it. During a severe deficit, operators know what must remain powered and what can disappear. The clarity reduces hesitation and prevents different teams from protecting their own equipment at the expense of system survival.

Design dossier — include restart energy in storage calculations

A battery may be able to support refuge loads for many hours yet lack the power or energy needed to restart pumps, compressors or generation equipment afterward. Storage sizing should therefore reserve enough capacity for the recovery sequence as well as the steady degraded load. Motor inrush or heating requirements can make restart more demanding than normal operation.

This is especially important near the lower state-of-charge boundary. Consuming the final reserve to extend refuge time can leave the settlement unable to black-start even after the external problem clears. Recovery energy is part of emergency margin.

Design dossier — coordinate electrical and thermal reserves

Batteries, converters and generators all produce heat or have temperature limits. A power emergency can therefore be worsened if thermal-control equipment is also constrained. The survival load list should include enough cooling or heating to keep the electrical hardware

capable of delivering the assumed reserve. Otherwise the calculated kilowatt-hours may be physically inaccessible.

Conversely, waste heat from generation can sometimes support habitat thermal needs if the architecture is designed for it. Cross-system energy accounting can improve efficiency, but the settlement should still survive when one coupling is unavailable.

Design dossier — distinguish redundancy from shared infrastructure

Two generators feeding one converter and one cable route do not create two fully independent power paths. A fire, connector failure or software fault in the shared element can remove both. Reliability diagrams should identify common buses, cooling loops, control networks and fuel systems so the crew understands which failures remain common-cause.

Physical separation, different protection zones or local storage can reduce that coupling. The goal is not to eliminate every shared component, which may be impractical, but to make the consequences explicit and protect the highest-value loads from one plausible common failure.

Design dossier — manage forecast error with rolling horizons

Instead of deciding the entire next week from one solar forecast, operators can update dispatch as new production and load data arrive. A rolling horizon preserves reserve while allowing flexible loads to move when the outlook improves. This reduces the cost of being conservative without betting survival on a long-range prediction.

The process needs decision thresholds: when to start dispatchable generation, when to pause industrial work, and what state of charge must be restored before optional loads return. Repeated operational data can refine those thresholds over seasons.

Design dossier — test protection coordination after expansion

Adding a new building or large motor can change fault currents and the sequence in which protective devices operate. A breaker setting that once isolated a small branch may now trip an upstream bus, unnecessarily removing multiple habitats. Electrical expansion therefore includes review of protection settings and fault studies appropriate to the system.

Commissioning tests verify that a local fault remains local where intended. This is the electrical equivalent of pressure compartmentation: the network should contain damage rather than propagate it.

Design dossier — make energy governance auditable

When power is scarce, different teams will have legitimate reasons to protect laboratories, agriculture, manufacturing or comfort. A transparent priority framework prevents informal influence from replacing system logic. The control room records why a priority was changed, who authorized it and what margin remained.

This record also improves planning. If the same load repeatedly receives emergency protection, perhaps its normal classification is wrong or local backup should be added. Governance data can expose design needs just as telemetry does.

Final synthesis for mission qualification

Final synthesis — power security is preserved margin, not installed megawatts

A power-secure settlement demonstrates that essential loads remain supplied through the selected source outage, storage degradation and distribution fault while still retaining a path to black start. Installed generation matters, but so do common buses, cooling dependencies, startup energy and the discipline to keep reserve unallocated during normal growth. Reviewers should be able to trace each critical load to at least one credible supply path and identify the exact stage at which flexible loads shed if margin deteriorates.

Final synthesis — operations should make the invisible reserve visible

Dashboards and daily planning need to show more than present kilowatts. They should show usable stored energy, forecast production range, largest credible source loss, protected black-start energy and the routine load that can be removed at each shedding stage. When these margins are visible, crews can act before an emergency threshold is crossed. When they are hidden inside engineering spreadsheets, the settlement can appear healthy until one ordinary maintenance outage reveals that all spare capacity has already been consumed.

Qualification notes

Qualification note — verify reserve after each major configuration change

Replacing a battery, adding a habitat or changing generation-control software can alter the effective reserve even when installed nameplate capacity increases. After major changes, the team should repeat peak-load, daily-energy, shedding and black-start checks with the new configuration. Commissioning evidence should include measured startup behavior and the new protected reserve rather than assuming previous qualification still applies.

Qualification note — keep restoration priority separate from normal load priority

The order in which loads return after a grid disturbance may differ from their normal criticality. Some support equipment must be energized before the critical function it serves can restart, while large critical loads may need to wait until generation stabilizes. A restoration table should therefore show dependencies and sequencing, not simply reverse the shedding list.

Final margin note

A resilient microgrid should still show a positive recovery margin after realistic conversion losses, delayed repair and forecast error are included. If the margin exists only in the ideal schedule, the system is operating too close to its boundary for long-duration settlement use.

Interactive beginner glossary

The energy vocabulary below is written as operational definitions, because a crew must connect each word to a decision it may need to make under pressure.

- power — The rate at which electrical energy is produced or consumed at a given moment, commonly expressed in kilowatts.
- energy — The accumulated electrical work over time, commonly expressed in kilowatt-hours.
- peak load — The highest relevant simultaneous electrical demand during the scenario being considered.
- load profile — A timeline showing how electrical demand changes over a period.
- base load — Demand that remains present for most or all of the operating period.

- flexible load — A load whose operating time can be shifted without unacceptable consequence.
- critical load — A load whose loss can quickly threaten safety, mission survival or an irreplaceable function.
- microgrid — A local electrical network that coordinates generation, storage, distribution and loads and can often operate independently.
- dispatchable source — A generator whose output can be commanded when needed within its operating limits.
- variable generation — Electrical production that changes with environmental conditions rather than following demand directly.
- state of charge — The fraction of a battery's available stored energy at a given time.
- usable capacity — The portion of nominal storage that operators allow for the current operating purpose.
- reserve margin — Capacity intentionally kept available to absorb uncertainty, faults or unexpected demand.
- load shedding — Deliberate disconnection or reduction of selected loads to protect the remaining electrical system.
- black start — Restarting generation and the grid when the main electrical system is de-energized.
- seed source — A protected small source that provides the initial energy needed to begin black start.
- islanding — Separating part of a microgrid so it can operate independently from a damaged or unstable section.
- bus — A common electrical distribution connection serving multiple sources or loads.
- feeder — A circuit carrying electrical power from a distribution point toward a load group or subsystem.
- converter — Equipment that changes electrical voltage, current form or other power characteristics.
- inverter — A converter that produces alternating current from a direct-current source.
- distribution loss — Electrical energy converted to heat or otherwise lost while moving and conditioning power.
- coincident load — Two or more loads demanding power during the same time interval.
- duty cycle — The fraction of time equipment operates within a repeating period.

- spinning reserve — Immediately available generation margin on running machines; the broader idea is rapid response reserve, even when the technology is not literally rotating.
- ride-through — The ability of a load or local source to continue operating through a short disturbance.
- synchronization — The checks required before connecting two energized electrical sections so they can operate together safely.
- protection coordination — Setting electrical protective devices so the smallest necessary section disconnects when a fault occurs.
- contingency — A credible abnormal event included deliberately in planning and capacity checks.
- dispatch — The operational choice of which sources, storage devices and loads should change output or consumption at a given time.

Operational review checklist

- Maintain separate peak-power and energy budgets.
- Classify each major load by consequence, safe interruption time and restart behavior.
- Build and update a measured load profile rather than relying on one average power number.
- Use credible solar or variable-generation envelopes rather than nameplate output alone.
- Account for maintenance and startup limits of dispatchable generation.
- Track battery nominal, current and intentionally reserved capacity separately.
- Include converter, cable and charge/discharge losses at clearly defined boundaries.
- Preserve an explicit reserve margin during normal operation and expansion.
- Define staged shedding with estimated power recovered at each stage.
- Test that automatic shedding remains understandable and manually recoverable.
- Protect a seed source and rehearse black start as a dependency sequence.
- Identify which habitats or functions can island after a distribution fault.
- Control reconnection so large loads do not collapse a weak island.
- Approve maintenance only after testing simultaneous demand and a second credible fault.
- Require capacity upgrades before routine growth consumes the contingency margin.

Microgrid calculation studio: power, energy, storage and survivable margins

A microgrid course must separate power from energy. Power is the instantaneous rate at which energy is transferred; energy is the accumulated quantity over time. Confusing kilowatts with kilowatt-hours makes otherwise sophisticated settlement studies unusable.

Formula A – energy consumed by a steady load

Question. How much electrical energy does a load consume while it operates?

$$E = P \times t$$

Read aloud. “Energy equals power multiplied by time.”

- **E:** energy, commonly kWh in settlement operations.
- **P:** electrical power, kW.
- **t:** operating time, h.

Teaching calculation. A 12 kW pump operates for 3 h. $E = 12 \times 3 = 36$ kWh.

Unit check. $\text{kW} \times \text{h} = \text{kWh}$. A kilowatt-hour is energy, not power.

Mental check. At 10 kW for 3 h the answer would be 30 kWh, so 36 kWh is plausible.

Limit. This assumes constant power. Variable loads require summing intervals or integrating power over time.

Formula B – instantaneous generation surplus or deficit

Question. At one moment, is generation exceeding demand or falling short?

$$P_{\text{net}} = P_{\text{generation}} - P_{\text{load}}$$

Read aloud. “Net power equals generation power minus load power.”

Teaching calculation. If available generation is 210 kW and load is 185 kW, $P_{\text{net}} = 25$ kW. If load rises to 230 kW, $P_{\text{net}} = -20$ kW.

Interpretation. Positive net power can charge storage or support optional loads. Negative net power must be covered by storage, another generator, import from another section, or load shedding.

Independent check. $P_{\text{generation}} = P_{\text{load}} + P_{\text{net}}$.

Limit. The scalar balance does not prove that cables, converters or local branches can carry the required power.

Formula C – usable stored energy

Question. How much of a nominal battery energy rating is intentionally available to operations after depth-of-discharge and conversion allowances?

$$E_{\text{usable}} = E_{\text{nom}} \times f_{\text{window}} \times \eta$$

Read aloud. “Usable energy equals nominal energy multiplied by the allowed operating-window fraction and by efficiency.”

- **E_{nom} :** nominal stored energy.
- **f_{window} :** fraction of the nominal capacity permitted by the operating policy.
- **η (“eta”):** modeled discharge/conversion efficiency.

Teaching calculation. With 500 kWh nominal, $f_{\text{window}} = 0.80$ and $\eta = 0.92$: first $500 \times 0.80 = 400$ kWh in the allowed window; then $400 \times 0.92 = 368$ kWh usable in this simplified model.

Unit check. $\text{kWh} \times \text{dimensionless} \times \text{dimensionless} = \text{kWh}$.

Limit. Real available energy also depends on temperature, age, rate capability, protection limits and the battery system actually selected.

Formula D – ideal critical-load endurance

Question. How long can a usable energy reserve support a steady protected load?

$$t_{\text{endurance}} = E_{\text{usable}} / P_{\text{critical}}$$

Read aloud. “Endurance time equals usable energy divided by critical power.”

Teaching calculation. With $E_{\text{usable}} = 368$ kWh and $P_{\text{critical}} = 92$ kW, $t = 368/92 = 4.0$ h.

Unit check. $\text{kWh}/\text{kW} = \text{h}$.

Reverse check. $92 \text{ kW} \times 4 \text{ h} = 368 \text{ kWh}$.

Operational meaning. Four hours is the ideal planning time at that load. It should be reduced by any reserved recovery energy and revised if the load changes.

Formula E – delivered-power margin

Question. How much deliverable power remains above the modeled peak protected demand?

$$M_P = (P_{\text{available}} - P_{\text{peak}}) / P_{\text{peak}}$$

Read aloud. “Power margin equals available power minus peak demand, divided by peak demand.”

Teaching calculation. If available deliverable power is 240 kW and protected peak is 200 kW, difference = 40 kW; $40/200 = 0.20 = 20\%$.

Mental check. The spare 40 kW is one fifth of 200 kW.

Limit. A 20% generation margin is not automatically a 20% system margin. A single converter, bus or cable can still be the limiting element.

Formula F – first-pass solar-array output

Question. In a deliberately simplified model, how do illuminated area, incident irradiance and conversion efficiency combine?

$$P_{\text{array}} = G \times A \times \eta$$

Read aloud. “Array power equals irradiance multiplied by active area and conversion efficiency.”

- **G:** incident solar power per square metre, W/m^2 .
- **A:** active collection area, m^2 .
- **η :** conversion efficiency as a decimal.

Teaching calculation. Use illustrative values $G = 450 \text{ W/m}^2$, $A = 100 \text{ m}^2$, $\eta = 0.25$. Incident power = $450 \times 100 = 45,000 \text{ W}$. Electrical output in this simple point model = $45,000 \times 0.25 = 11,250 \text{ W} = 11.25 \text{ kW}$.

Limit. This is not a Mars array design. Angle, dust, temperature, spectrum, wiring, degradation, storage and seasonal/diurnal variation must be modeled separately.

Formula G – protected-load fraction

Question. During a contingency, what fraction of the nominal load has been classified as protected?

$$f_{\text{protected}} = P_{\text{protected}} / P_{\text{nominal}}$$

Teaching calculation. If nominal settlement demand is 250 kW and emergency protected loads total 100 kW, $f_{\text{protected}} = 100/250 = 0.40 = 40\%$.

Interpretation. This does not say that 60% is unimportant. It shows the magnitude of deliberate load shedding required by the emergency state.

Design question. For each shed load, state the maximum time it may remain off before it creates a new hazard or recovery problem.

Integrated blackout exercise

A teaching microgrid has 600 kWh nominal storage. Policy allows 80% of nominal capacity and the exercise assumes 90% conversion efficiency. Protected loads are 90 kW. Calculate usable energy and ideal endurance. Then reserve 80 kWh for restart and diagnosis and recompute operational endurance.

Solution. Usable energy = $600 \times 0.80 \times 0.90 = 432$ kWh. Ideal endurance = $432/90 = 4.8$ h. If 80 kWh is protected for recovery, energy available for holding loads is 352 kWh; endurance = $352/90 \approx 3.91$ h. The difference demonstrates why a recovery reserve must be separated from “energy still in the battery.”

Microgrid calculation laboratory: margin and endurance are different questions

Power architecture must answer two separate questions: can generation meet the load now, and how long can stored energy protect critical functions when generation is insufficient?

Protected-load margin

Question. How much power margin exists above the protected load?

$$M = (P_{\text{available}} - P_{\text{protected}}) / P_{\text{protected}}$$

Read aloud. Margin equals available power minus protected power, divided by protected power.

- **P_{available}**: power that can actually be delivered now, in kW.
- **P_{protected}**: load that operations have decided must remain supplied, in kW.
- **M**: dimensionless margin; multiply by 100 for percent.

Teaching scenario. If 140 kW is available and protected loads total 110 kW, the difference is 30 kW. Dividing 30 by 110 gives 0.273, or 27.3% margin. If one generator derates, the margin can disappear without the protected loads changing.

Limit. This ratio says nothing about voltage stability, frequency control, startup transients, distribution faults or whether generation and load are geographically connected.

Battery endurance

$$t = E_{\text{usable}} / P_{\text{critical}}$$

If a teaching battery bank stores 320 kWh but operations preserve 20% as an unavailable protection margin, usable energy is $320 \times 0.80 = 256$ kWh. With an 80 kW critical load, ideal endurance is $256/80 = 3.2$ h. Real endurance must include conversion losses, temperature, degradation and power limits.

Unit check. kWh $\div$ kW = h.

► Exercise — degraded generation

First-Man power desk: keep frequency, voltage and life-support priorities coherent during a fault

A settlement microgrid is not just generation plus batteries. It is a control system deciding, every second and every hour, which sources connect, which loads remain energised and how stored energy is protected. The design should classify loads by consequence and restart behaviour. Some loads are continuously vital, some can pause for minutes, others can be deferred for hours, and some industrial loads are valuable precisely because they can be scheduled when power is abundant.

Sense

generation, bus state,
storage, forecasts

Protect

isolate fault, preserve
stable bus

Prioritise

life support before
deferrable work

Recover

black start, staged
reconnection, verify
margin

Use the power/energy distinction during fault dispatch

This is an application of the distinction established in the opening lesson, not a second definition. During a fault, operators first ask whether available generation can carry the protected loads now; they then ask how long stored or dispatchable energy can sustain that configuration. A plan that answers only one question can still fail the crew.

Black-start capability deserves explicit testing. If the settlement loses the main bus, what independent source can energise controls, communications, valves and one generation path? Which loads reconnect automatically, and which must remain inhibited until the bus is stable? A recovery plan that assumes all loads can return simultaneously can trip the system again because restart peaks are often larger than normal steady demand.

Forecast uncertainty belongs in the dispatch decision

Solar generation, dust loading and maintenance availability are uncertain. Operators should not spend storage down to the nominal overnight requirement if the morning generation forecast is fragile. A reserve policy should identify the consequence of a forecast miss and protect the battery state needed to survive it. The more uncertain the incoming generation, the more valuable flexible industrial loads and thermal storage become.

Exercise: choose what to shed first

A fault removes one generator while the battery is at 42% state of charge. The current load is 85 kW, but protected generation plus allowable battery discharge can sustain only 62 kW without violating reserve. Life support, medical storage, communications and essential thermal control total 49 kW. A 14 kW chemical batch and 9 kW science load are running. The first action is to stop or defer the 23 kW of nonessential work, bringing demand to 62 kW. The important second action is to decide whether 62 kW itself is sustainable over the expected repair time; instantaneous balance is not the same as endurance.

Black start is a dependency sequence: control and protection must exist before generation can safely energise critical loads. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Operational qualification lab: keep a microgrid alive after the easy reserves are gone

A settlement power budget becomes operationally meaningful only when it survives a source loss, a bad forecast and a load that refuses to behave exactly as planned. Nameplate generation is therefore less useful than **verified available power** under a stated condition. The crew must know which loads are protected, which can be deferred, which can be throttled and which create a larger hazard if power is removed abruptly.

Build the load ladder before the emergency

Classify loads by consequence rather than by convenience. Atmosphere circulation, critical thermal control, communications needed for emergency coordination and medical support may sit above workshop production or discretionary science. Some loads are not binary: a greenhouse may tolerate reduced lighting for a period; a pump may have a lower safe operating point; a chemical process may need an orderly shutdown rather than immediate disconnection. A credible priority table therefore records nominal power, minimum protected power, maximum interruption, restart surge and shutdown constraints.

Microgrid topology matters as much as aggregate kilowatts. Two generators on one vulnerable bus are not equivalent to two sources that can feed protected loads through independent paths. A contingency review must ask what remains routable after the failed component is isolated.

Protected-load power margin after a contingency

$$M_P = (P_{\text{available,verified}} - P_{\text{protected}}) / P_{\text{protected}}$$

1 – Concrete question

After a defined generator, feeder or storage failure, how much verified power remains above the protected load?

2 – Intuition

Compare what the surviving architecture can truly deliver with what the settlement must keep powered.

3 – Quantities

Use verified deliverable power after derating and routing limits, and the protected load after the planned shedding sequence.

4 – Formula

Margin is surviving verified power minus protected demand, divided by protected demand.

5 – Read aloud

"M P equals verified available power minus protected power, divided by protected power."

6 – Symbols

M_P is dimensionless margin; $P_{\text{available,verified}}$ is contingency deliverable power; $P_{\text{protected}}$ is the power that must remain supplied.

7 – Pronunciation

Subscripts describe the operating state; they are not multiplication.

8 – Units

Both powers use watts or kilowatts. The ratio is dimensionless and can be expressed as a percentage.

9 – Convention

Define whether storage discharge is included and for what duration. A battery that can sustain 100 kW for ten minutes is not 100 kW of indefinite generation.

10 – Why this relationship

The numerator is spare power; dividing by protected demand makes the spare comparable across different settlement sizes.

11 – Assumptions

Frequency, voltage, thermal ratings and routing are assumed acceptable at the stated operating point.

12 – Unit check

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

13 – Numerical case

Verified surviving delivery: $P_{\text{available,verified}} = 510$ kW.

Protected demand: $P_{\text{protected}} = 425$ kW.

Surplus = $510 - 425 = 85$ kW.

$$M_P = 85 / 425 = 0.20.$$

$$M_P = 20\%.$$

14 – Operations

Subtract to find the 70 kW spare, then divide by 350 kW to express that spare relative to the protected demand.

15 – Algebra check

$$P_{\text{available}} = P_{\text{protected}}(1+M). \quad 350 \times 1.20 = 420 \text{ kW}.$$

16 – Mental estimate

Seventy is one fifth of 350, so a 20% margin is plausible.

17 – Interpretation

The settlement can absorb a modest demand increase, but the margin is not large enough to ignore a forecast error or restart surge.

18 – What it does not prove

It does not prove energy endurance, fault-clearing performance, cable capacity or stability of the microgrid.

19 – Sensitivity

If protected demand rises to 390 kW, margin falls to about 7.7%. Small load-growth errors can therefore consume the contingency margin quickly.

20 – Practice

Guided exercise. Calculate the protected-load margin for 510 kW verified available power and 425 kW protected demand.

Detailed guided correction.

1. Surplus = $510 - 425 = 85$ kW.
2. Divide surplus by protected demand: $85 \div 425 = 0.20$.
3. Margin = 20%. A 20% margin is not the same as 20% battery endurance; energy duration must be checked separately.

Autonomous exercise. The same settlement has a battery able to sustain a 90 kW deficit for 30 minutes. A generator trip reduces verified generation to 360 kW while protected loads are 425 kW. Determine the immediate deficit, battery energy consumed if the condition lasts 20 minutes, and a defensible shedding decision.

► Autonomous correction – open after attempting the exercise

21 – Mission decision

Set a HOLD threshold for adding new loads whenever the N-1 margin or endurance falls below the settlement's accepted

criterion.

Energy is the second clock

Power answers whether the lights can remain on now; energy answers for how long. During a dust event or generator outage, the crew should track both the instantaneous power margin and the energy reserve. A technically “positive” power margin can still hide a battery depletion deadline. Conversely, a large stored-energy inventory is useless if the inverter or feeder cannot deliver the required peak power.

Qualification drill

Create a load-shedding table for twelve loads. Include one load that must be ramped down, one with a high restart surge, one whose interruption is acceptable only for two hours and one that is normally low priority but becomes safety-critical during an EVA. Inject a feeder loss and a forecast error. Your answer should specify the switching order, the new protected load, the battery clock and the evidence required before restoring deferred loads.

Source context. NASA power technology material and the NASA Systems Engineering Handbook are used here as primary engineering context; the contingency numbers are teaching assumptions. [NASA Systems Engineering Handbook](#).

R61 microgrid restoration doctrine: black-start, protection, dispatch and operator evidence

A settlement microgrid can have enough generators and storage on paper yet still fail to recover from a blackout. Recovery requires a sequence: establish an electrically stable island, energise control and communications, start the resources capable of black-start, restore life-critical buses, verify frequency and voltage, then add loads without exceeding transient or energy limits. The course therefore needs restoration logic as well as nominal power arithmetic.

Identify the black-start path before the first blackout

Not every generator can start a dead grid. Some need auxiliary power, thermal conditioning, communications or excitation. The restoration plan should name the source that can energise essential controls from zero, the bus it first powers, the loads that must remain disconnected,

and the conditions for connecting the next source. If the designated black-start battery also feeds emergency lighting and communications, its energy budget must reserve enough capacity for the restoration sequence.

Primary-source bridge — NASA power systems. NASA's power-system work provides primary context for generation, storage and power management. R61 extends that logic to a surface microgrid restoration sequence with protected loads and explicit black-start dependencies. [Official source](#).

Protection coordination matters when the grid changes shape

A breaker or software trip threshold that is appropriate with all generators online may behave differently when the settlement is operating on one inverter and a battery. Fault current, inrush and voltage recovery can change. The team should therefore qualify protection in representative islanded states, not merely assume settings proven in the full grid remain selective. A protection system that trips the healthy backup during a fault can turn redundancy into a common-cause outage.

Restoration priority is not identical to normal operational priority

Some loads are valuable in steady operation but harmful during restart because they have high inrush, unstable control behaviour or large thermal transients. Build a restoration ladder that first recovers the loads needed to keep people alive and to observe the grid, then restores environmental control, communications, water and thermal systems in a sequence compatible with generation and storage. Industry and discretionary loads follow only after margin is demonstrated.

Track both instantaneous margin and remaining energy

A battery can supply a high protected load for minutes while the energy state of charge is collapsing. Conversely, a generator may have enough energy potential but insufficient instantaneous power to start a motor. The dispatch board should show power margin, stored-energy endurance, expected generation, uncertainty and the next decision point. These quantities answer different questions and should never be collapsed into a single green status light.

Primary-source bridge — Systems Engineering Handbook. NASA systems-engineering guidance is relevant to interfaces and verification. A microgrid restoration plan is an interface-rich sequence that must be verified as an integrated behaviour, not inferred from component nameplates. [Official source](#).

Control-system common cause must be tested

Multiple generators controlled by the same network switch, time source, software build or supervisory controller may not be independent. Include at least one exercise in which the normal supervisory layer is unavailable. The crew should be able to identify which local controls remain, which protections are autonomous, and how much of the grid can be recovered safely without creating uncontrolled parallel sources.

Restoration exercise — stable power exists, but the wrong load is next

After a blackout, one battery inverter and one generator have stabilised a protected bus. Oxygen production is down but reserve oxygen is adequate for several hours. A large water processor can restart immediately, while the thermal pump serving the active inverter has only twenty minutes of thermal margin. The correct priority is the thermal pump even if water appears higher in the normal mission hierarchy. The exercise teaches that restoration sequence follows time-to-consequence and system dependencies, not a static list.

Black-start rehearsals should include failed starts

A restoration script that succeeds only when every step works on the first attempt is not robust. During training, inject a generator that fails to synchronise, a stale breaker-status indication or a battery that reaches its lower state-of-charge limit earlier than predicted. The crew should know the branching logic: when to retry, when to abandon a source, when to preserve remaining energy and when to reduce the protected load list.

Power-quality evidence belongs beside megawatts and kilowatt-hours

Sensitive avionics, chargers and variable-speed drives can fail or trip even when average voltage appears acceptable. Where the architecture requires it, trend voltage excursions, frequency behaviour, harmonics or transient events during switching. A grid that delivers the nominal energy but repeatedly destabilises critical electronics is not electrically healthy.

Restoration should end with a configuration reconciliation

Emergency switching can leave tie breakers, local overrides, temporary generator limits or disabled automatic functions in unusual states. Before returning to normal dispatch, compare the actual electrical one-line and controller modes with the documented configuration.

Recovery is incomplete if the next crew inherits a grid that works but is no longer understood.

Operators must protect both the instantaneous power margin and the remaining energy clock throughout restoration. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.