

Operational governance: decisions, tradeoffs and crisis command



Conceptual visualization — governance becomes operational when authority, evidence, time limits, dissent, handover and resource priorities meet in one decision process.

Define roles, authority, escalation and crisis decisions when Earth cannot command a Mars settlement in real time.

Before starting — reconnect mission operations, delayed communications and procedure design before assigning authority during a settlement crisis. Important quantities and assumptions are stated at first use.

Mastery objectives

- separate settlement governance, operational command, technical authority and emergency delegation
- classify decisions by reversibility, time margin, consequence and the information needed before commitment
- design escalation and dissent recording that still works with Earth-Mars communication delay
- restore normal authority after a crisis without losing the evidence needed for learning and accountability

1. Time delay changes governance

A Mars settlement cannot ask Earth for permission during every incident. Communications delay forces delegation while common rules and mission objectives still matter. Operational governance defines what the crew may decide locally, what must be reported and what requires delayed concurrence.

2. Role, authority and accountability


A useful organization chart does more than name positions. It specifies authorized decisions, required information, backup personnel and limits. A power lead may shed industrial loads but should not disable a critical medical barrier without coordination. Clear authority prevents a technical crisis from turning into a jurisdiction dispute.

3. Normal, degraded and emergency states

Decision rules change with system state. Normal operations favor planning and review. In degraded mode, local leads may act more quickly within a defined envelope. In an emergency, immediate authority protects life and habitat integrity. Transitions between states should be declared and logged.

Crisis authority map

A decision timeline separates technical advice, incident authority, crew consequence and command commitment



Rule: authority must arrive before the system crosses an irreversible physical state.

DELTA-SIERRA · SPACE ACADEMY · R63 ATLAS CLOSURE

A delayed Mars operation needs clear boundaries between expertise, incident command, governance and remote advice. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

4. Reversible and irreversible decisions

Crisis discipline distinguishes actions that can be tried and reversed from actions that permanently consume a resource or close an option. Shedding a noncritical load is reversible; venting a rare reagent is not. The more irreversible the action, the more carefully assumptions, sensors and alternatives should be checked unless life is immediately threatened.

5. Technical dissent

Specialists can interpret the same data differently. Procedure should preserve the right to raise dissent without paralyzing action. Teams state assumptions, consequences and evidence that would change the decision. An authority then decides, and the technical dissent remains recorded for later review.

6. Priority rules

The settlement publishes priorities before a crisis: human life, refuge integrity, atmosphere control, communications, preservation of scarce resources, then scientific and industrial objectives. These priorities do not replace judgment, but they reduce improvised bargaining when several systems compete for the same power or personnel.

7. Decision log

Important decisions record time, known state, available data, options considered, authority, action and expected result. This is not paperwork for its own sake. It allows the next shift to understand unusual configurations and supports post-incident analysis. An undocumented workaround can become a delayed failure.

8. Relationship with Earth

Ground teams remain valuable for analysis, expertise and logistics. Messages should be structured for delay: state, timeline, precise question, required data and actions already taken locally. Earth should not receive a large telemetry dump without knowing what decision the crew is trying to support.

Deepening: command handover

A crisis can last for days. Command must be transferable. Handover includes system state, hypotheses, active work, upcoming decisions and critical resources. Fatigue of the decision-maker is itself an operational hazard; a design that requires one irreplaceable person to remain awake for thirty hours is a human-system failure.

Deepening: minority expertise

Complex decisions may place majority opinion against specialist expertise. Robust governance defines domains where the specialist position must be explicitly addressed before proceeding, such as radiation exposure or pressure safety. The expert does not control every decision, but overriding specialist advice should be conscious, justified and recorded.

9. Worked example: decision cycle with delay

With a 14-minute one-way delay, a question sent to Earth and an immediate reply require at least 28 minutes of radio travel. If the ground team needs 20 minutes to analyze the data, the minimum cycle is $14 + 20 + 14 = 48$ minutes. Any action required within ten minutes cannot depend on that loop.

Calculated case study: can a crisis decision wait for Earth?

TEACHING ASSUMPTION — One-way radio delay is 12 min. Combined ground preparation, reading and validation add 6 min. A local safety decision must be taken within 10 min.

Let t_{one} be one-way delay in minutes; t_{ground} total ground processing in minutes; t_{round} the minimum decision round trip; and t_{window} the local decision window.

$t_{\text{round}} = 2 \times t_{\text{one}} + t_{\text{ground}} = 2 \times 12 + 6 = 30$ min. Difference from the window: $\Delta t = t_{\text{round}} - t_{\text{window}} = 30 - 10 = 20$ min. An Earth response would therefore arrive 20 min after the assumed window closed.

The calculation does not decide who should hold authority. It shows why Mars governance must predefine delegated authority and local rules for emergencies incompatible with communication delay.

10. Exercise

A power failure removes 35% of capacity while agriculture, workshop, laboratory and medical systems all request power. Build a priority matrix, define who can shed loads and specify what should be sent to Earth after stabilization.

11. Reasoned solution

The power lead applies pre-agreed rules: life-support and safety loads are protected first, deferrable activities are reduced. Medical systems are not automatically run at maximum if no intensive procedure is underway, but critical medical functions remain protected. The decision is reevaluated as fault diagnosis improves and is logged for handover.

12. Mini-project

Write an operational charter for a 30-person settlement: normal/degraded/emergency states, six functional authorities, delegation rules, dissent process, decision log and Earth communication protocol. Test it against a combined casualty, power-loss and intermittent-communications scenario.

Operational governance is an engineering interface between authority, information and time

A Mars settlement cannot treat governance as a document that matters only when people disagree. In a remote settlement, governance determines whether a technical warning reaches the person who can act, whether authority transfers when a leader is unavailable, whether scarce resources are allocated consistently and whether emergency decisions remain reviewable

afterward. The design problem is therefore operational: define who may decide, with which information, inside which limits, and how the organization changes state when time, data or people are constrained.

Earth remains an essential analytical partner, but physics prevents it from acting as a real-time supervisor. Depending on planetary geometry, one-way radio delay can be many minutes. A message may also wait for processing, analysis and validation before the reply is transmitted. The settlement therefore needs delegated authority that is explicit before a crisis. Delegation is not independence from Earth; it is a way to preserve common goals when the communication loop is slower than the decision window.

The most robust model separates three layers. **Political or institutional governance** defines long-term rules, rights and objectives. **Operational command** coordinates people and resources in real time. **Technical authority** protects domain-specific safety limits such as pressure integrity, radiation exposure, medical practice or electrical protection. These layers can cooperate without being identical. A commander can set priorities, but should not silently redefine an engineering safety limit during a routine scheduling dispute.

Decision rights and delegation

Every critical function should have a named primary authority, a backup and a documented scope. The power lead may be authorized to shed industrial loads when generation falls below a threshold. The medical lead may reserve power or oxygen for a clinical intervention. The habitat lead may isolate a compartment if pressure integrity is uncertain. The important point is not the title; it is the boundary of authority and the conditions under which that authority expands or contracts.

Delegation should be written as a decision envelope rather than a vague instruction to “use judgment.” An envelope can include the system state, the maximum resource that may be committed, the actions that require a second competent person, the actions that are prohibited without broader concurrence and the reporting requirement after the event. This makes delegation auditable and easier to rehearse.

Backup authority must be credible. A deputy who has never seen the decision tools or who lacks access to the logs is not a real backup. Handover drills should therefore be treated like technical redundancy tests: remove the nominal decision-maker and verify that the organization can continue safely.

Normal, degraded, emergency and recovery states

One set of rules cannot serve every condition. In normal operations, planning, peer review and scheduled coordination can dominate. In degraded mode, the settlement has lost margin but still controls the situation; local leads may need wider authority and faster coordination. In emergency mode, immediate protection of life and habitat integrity takes priority. Recovery mode begins only after the event is stabilized and focuses on rebuilding barriers without recreating the original failure.

State transitions should use measurable criteria. “Things look bad” is not a robust trigger. A transition can instead depend on loss of a redundant power channel, refuge occupancy, atmosphere parameters outside defined limits, communications unavailable beyond a stated duration, or remaining water reserve below a threshold. The thresholds themselves require engineering validation, but the governance principle is general: the organization should know when its rules change.

Recovery deserves its own state because emergency workarounds create unusual configurations. Temporary cables, bypassed automation, relocated medical stock and personnel working outside normal shifts can remain hidden after the immediate event. Recovery governance assigns ownership to each temporary condition and requires an explicit decision before the settlement returns to nominal status.

Priority rules for scarce resources

A crisis becomes difficult when two legitimate functions need the same constrained resource. Power, cooling, communications bandwidth, technician time, rover access and protected consumables can all become bottlenecks. Priority rules should be established before the crisis and should focus on consequences rather than departmental status.

A useful hierarchy starts with preservation of life, pressure and atmosphere integrity, fire control, medical stabilization and refuge capability. Communications and diagnostic capability usually follow because they support the recovery of other systems. Food production, research and industry can be essential over longer timescales but may be temporarily reduced. The exact hierarchy depends on the scenario: a greenhouse cannot be ignored for weeks, and a workshop may become critical if it is repairing the only failed life-support component.

For that reason, priorities should be dynamic. The decision team should identify the *time to consequence* for every competing load. A medical ventilator may have an immediate consequence, a water purification unit may have hours or days of buffer, and a fabrication job may have several days. This converts a political-looking conflict into a time-and-consequence problem that can be reasoned about transparently.

Decision logs: the settlement's operational memory

A strong decision log records the time, system state, known facts, uncertain facts, assumptions, options considered, authority, action, expected effect, monitoring plan and next review point. It should also preserve technical dissent when specialists disagree. The purpose is not blame. The log lets the next shift understand why the settlement is in an unusual configuration and allows later analysis to separate a bad outcome from a bad decision made with the information available at the time.

Logs should distinguish observations from interpretations. “Tank pressure sensor A reads 18% below sensor B” is an observation. “Sensor A has failed” is a hypothesis until corroborated. Mixing the two is dangerous because later decisions may treat an untested interpretation as fact. A structured log can therefore include confidence levels or evidence status alongside each critical statement.

Decision logs are also part of knowledge transfer. A settlement that survives many anomalies but does not capture why its workarounds succeeded will repeatedly rediscover the same lessons. Post-event reviews should feed procedures, training scenarios, spare-part policy and software changes.

Technical dissent without paralysis

Complex systems create legitimate disagreement. A medical specialist may favor one risk trade while a power engineer sees a different system consequence. Governance should guarantee that dissent can be raised, recorded and answered without requiring unanimous agreement. The decision authority still has to decide within the available time.

A useful dissent format is concise: state the concern, the evidence, the predicted consequence, the recommended alternative and the evidence that would change the specialist's view. This makes dissent actionable. It also prevents the discussion from degrading into authority-by-volume or status.

Some domains should carry protected technical authority. For example, a pressure vessel, electrical protection, radiation exposure or medical procedure may have certified limits that operational convenience cannot override casually. If an emergency forces departure from a limit, the deviation should be explicit, time-bounded where possible and documented as an exceptional risk decision.

Earth as delayed partner, not remote joystick

Messages to Earth should be designed around a decision request. A compact crisis message can include current state, timeline, confirmed measurements, uncertain points, actions already taken, resources remaining, alternatives under consideration and the exact question for ground specialists. Sending every raw telemetry channel without context can slow analysis rather than improve it.

Earth can provide deep expertise, large computational resources and independent review, but it should not be placed on the critical path of actions that must occur before a reply can arrive. Procedures should state which actions are local by default and which strategic changes wait for Earth when time allows.

When communication is intermittent, the settlement can send decision packages that remain useful even if the reply is delayed. Ground teams can likewise return conditional guidance: “if pressure trend X remains below Y, use option A; otherwise use option B.” This makes the communication loop more resilient to delay and missing packets.

Calculation laboratory: delay, deadlines, workload and resource arbitration

Minimum Earth decision cycle

$$t_{\text{cycle}} = 2 \times t_{\text{one_way}} + t_{\text{ground}} + t_{\text{local_processing}}$$

Read aloud: the minimum decision cycle equals twice the one-way light-time, plus ground analysis time, plus local time needed to read and validate the response. If one-way delay is 13 minutes, ground analysis takes 18 minutes and the local team needs 4 minutes to interpret the response, the minimum cycle is $2 \times 13 + 18 + 4 = 48$ minutes. A decision due in 12 minutes cannot rely on that loop.

Deadline margin

$$M_{\text{time}} = t_{\text{deadline}} - t_{\text{cycle}}$$

If the deadline is 60 minutes and the cycle is 48 minutes, margin is 12 minutes. A positive margin does not automatically mean waiting is wise; uncertainty and additional communication rounds can consume the remaining time. A negative margin proves that the external loop is physically incompatible with the deadline.

Decision workload ratio

$$W = T_{\text{required}} / T_{\text{available}}$$

If a crisis cell requires an estimated 18 person-hours of analysis during the next two hours and has six people who can each contribute two effective hours, available capacity is 12 person-hours and the workload ratio is $18/12 = 1.5$. The team must simplify, delegate or add support because the planned work exceeds available attention.

Priority by time to consequence

$$\text{Urgency}_i = 1 / t_{\text{consequence},i}$$

This simplified teaching metric makes one idea visible: the shorter the time to serious consequence, the higher the urgency. It must not be used as an automatic decision algorithm because severity, reversibility and uncertainty also matter. It is a discussion aid for comparing demands that otherwise appear incomparable.

Protected reserve after allocation

$$R_{\text{after}} = R_{\text{before}} - \sum \text{allocation}_i$$

If 120 kWh of protected battery energy is available and emergency loads consume 25, 18 and 12 kWh during the planned interval, 65 kWh remains. The log should show both the allocation and the trigger for revisiting it. Consuming a protected reserve without a review point can turn one crisis into a later one.

Integrated crisis drill: two vital functions request the same power

Assume a generation fault leaves 70 kW of discretionary capacity after the habitat's protected baseline is served. A medical procedure requests 45 kW for one hour, while a water-processing recovery sequence requests 40 kW for ninety minutes. Both requests are legitimate, but they cannot run at full demand simultaneously.

The team first establishes time to consequence. The medical lead states whether delay changes patient risk and whether a reduced-power mode exists. The water lead states current tank inventory, contamination status and how long the recovery can wait. The power lead confirms the actual available margin and uncertainty in generation. The commander does not invent a technical answer; the commander integrates the consequences supplied by each authority.

If the medical procedure is time-critical and water storage provides several hours of buffer, the decision may allocate 45 kW to medicine and defer or sequence water recovery. The decision log records the remaining water margin and the exact time at which the deferral must be reassessed. If new information shows water quality deteriorating faster than expected, the decision is reopened.

Contradictory information drill

Suppose two oxygen sensors disagree after a ventilation anomaly. One indicates a safe concentration while another indicates a value outside the normal band. The governance failure would be to let the most senior person choose the preferred sensor. The technical response is to identify sensor independence, compare other evidence, check trends and move to a conservative state while uncertainty remains.

The decision log should preserve the disagreement rather than overwrite it with a single “official” number. If a third independent measurement becomes available, the team can update confidence. Governance here is inseparable from instrumentation architecture: independent authority is only useful if it has independent evidence.

Command-unavailable drill

A settlement should rehearse the sudden unavailability of its commander during a simultaneous technical incident. The deputy assumes authority using the same dashboards, procedures and communication channels. Functional leads continue within their delegated envelopes. Earth is informed, but the settlement does not wait for Earth to nominate a local replacement.

The drill passes only if everyone can identify the new authority, work does not pause for an improvised election or argument, and the handover log captures active risks. Afterwards, the team should examine whether too much knowledge or access had been concentrated in one person.

Progressive exercises with solutions

- ▶ Exercise 1 — Can Earth answer in time?
- ▶ Exercise 2 — Deadline margin
- ▶ Exercise 3 — Protected energy
- ▶ Exercise 4 — Dissent
- ▶ Exercise 5 — State transition

Operational charter mini-project

Produce a charter for a thirty-person settlement. Include at least six functional authorities, their backups, normal/degraded/emergency decision envelopes, protected technical limits, the command succession chain, the decision-log template, Earth communication format, dissent process and recovery-state exit criteria. Then test the charter against three events: two vital loads competing for power, contradictory atmosphere data, and the commander becoming unavailable.

The project should identify at least three decisions that must remain local because communication delay is too long, three decisions that can wait for Earth, and three decisions that require protected technical concurrence. The final document should make clear which values are engineering thresholds and which are governance choices.

Interactive beginner glossary

- decision authority — permission to decide within a defined scope.
- delegation — authorized transfer of decision power.
- degraded mode — controlled operation with reduced margin or capability.
- technical dissent — documented expert disagreement.
- time to consequence — remaining response time before serious impact.
- decision log — traceable operational memory of a decision.
- handover — structured transfer of command or work.
- recovery state — controlled return from emergency configuration.

Governance under pressure: separate technical authority, command authority and legitimacy

Operational governance becomes difficult when several forms of authority overlap. A commander may coordinate the whole settlement, while a physician controls a medical procedure, an electrical lead protects a power-system limit and a safety officer controls entry into a hazardous area. Good governance does not pretend these roles are identical. It defines where each authority begins, where it ends, and how conflicts are resolved when time is short.

Decision rights should follow consequence and expertise

Routine decisions can be delegated close to the work. High-consequence or irreversible decisions may require broader concurrence when time permits. Emergency decisions may temporarily concentrate authority, but the conditions for entering and leaving that emergency state should be explicit. Otherwise exceptional power can persist after the technical reason for it has disappeared.

The governance document should distinguish at least four things: who may propose an action, who supplies technical evidence, who may authorize it, and who verifies the result. Combining all four in one person creates speed but also creates blind spots and weakens independent checking.

Decision-chain capacity

$$C_{\text{decision}} = N_{\text{steps_completed_in_time}} / N_{\text{steps_required}}$$

This is not a score of legitimacy. It is a teaching indicator for whether the planned decision process physically fits the available time. If a process requires five essential steps—detection, technical assessment, authority decision, execution and verification—but only four can be completed before the consequence deadline, the process design is incomplete even if every step is individually sensible.

Information quality needs ownership

During an anomaly, data arrive with different reliability. A sensor may be calibrated or suspect, a crew observation may be direct or second-hand, a model may be validated or approximate, and a forecast may depend on uncertain assumptions. Decision logs should preserve those differences. Replacing all uncertainty with one authoritative number makes the record easier to read and easier to misuse.

A practical decision board can therefore separate confirmed observations, interpretations, assumptions and unresolved contradictions. The board should record what evidence would change the current decision. This prevents the team from defending yesterday's assumption merely because it has already been written down.

Dissent is a safety mechanism when structured

Technical dissent should be easy to raise and difficult to erase. The specialist states the disputed assumption, supporting evidence, predicted consequence, preferred alternative and the observation that would make the objection disappear. The decision authority then either accepts the recommendation or records why another option is chosen. This preserves both decisiveness and traceability.

Unstructured dissent, by contrast, can consume scarce time. The governance system should therefore teach concise formats before a crisis. Training matters because people under pressure revert to practiced patterns.

Command succession must include access and knowledge

Naming a deputy is not enough. The successor must have credentials, system access, current status information and authority recognized by the rest of the crew. A succession drill should test whether the settlement can continue when the commander is unavailable without waiting for Earth to resolve local authority.

The same principle applies to specialized roles. If only one person can authorize or understand a life-critical process, the settlement has a personnel single point of failure. Cross-training and documented delegation reduce that vulnerability without pretending every crew member has equal expertise.

Recovery governance closes temporary exceptions

Emergencies generate bypasses, temporary cables, altered software settings, rationing rules, relocated stocks and unusual shift patterns. Each temporary condition should have an owner, a review time and an exit criterion. Otherwise the settlement can declare “recovery complete” while hidden degraded configurations remain.

This is where operational governance and configuration management meet. A decision is not closed merely because the immediate danger has passed. The record, system configuration, reserves and procedures must again describe the state that physically exists.

Governance exercise: write three envelopes

For power, medicine and EVA, write a normal authority envelope, degraded envelope and emergency envelope. For each one, state who decides, which technical limits cannot be overridden without explicit exceptional action, what evidence must be recorded, how long the decision remains valid, and what restores the normal chain. Then introduce a communications outage and verify that the local process still functions.

Crisis-decision laboratory: quantify the time left to decide

Governance becomes operational when a decision has a consequence clock. Authority cannot create time that physics has already removed.

Decision-chain slack – first worked example

$$M_time = t_consequence - t_detect - t_diagnose - t_act$$

Teaching scenario. A condition is expected to become unacceptable in 90 min. Detection consumes 10 min, diagnosis 20 min and the selected action needs 15 min before it has effect. Remaining margin = $90 - 10 - 20 - 15 = 45$ min.

Interpretation. That 45 min is not “free time.” It contains uncertainty, communication, verification and possible failed attempts. A negative value means the proposed decision chain is physically too slow.

Backup coverage of critical authorities

$$C_backup = N_roles_with_trained_backup / N_critical_roles$$

If seven of eight critical operational roles have a trained designated backup, coverage is $7/8 = 87.5\%$. The missing eighth role is not “12.5% unsafe”; the ratio simply identifies an unresolved single-person dependency that deserves explicit treatment.

► Exercise — time margin

Conceptual visualisation: crisis governance should make authority, evidence, deadlines, dissent and handover visible before a decision becomes irreversible.

Decision architecture studio: who may decide, with what evidence, and before what deadline?

Governance becomes operational when authority is attached to specific decisions rather than to vague titles. A Mars settlement can distinguish at least four kinds of authority: technical authority over safety limits and engineering configuration; operational command over coordinated actions; medical authority over clinical decisions; and community or institutional authority over rules that affect rights, obligations and long-term priorities. In calm periods these authorities can overlap. In crisis they must be explicit enough that a crew does not waste its remaining margin negotiating who is allowed to act.

Build a decision-rights matrix

List recurring high-consequence decisions as rows: isolate a pressure zone, shed agricultural power, initiate a rescue sortie, quarantine a medical stock, enter refuge mode, override an automation, accept a repaired pressure vessel, ration a protected consumable. Columns identify who proposes, who verifies evidence, who authorizes, who must be informed and who can stop the action for a documented safety reason. This resembles a responsibility matrix, but it is built around decisions and consequences rather than administrative tasks.

Decision slack

$$S_{\text{decision}} = T_{\text{consequence}} - (T_{\text{detect}} + T_{\text{diagnose}} + T_{\text{coordinate}} + T_{\text{act}})$$

Question. How much time remains after the expected decision chain completes?

Units. Every T term is time, normally seconds, minutes or hours. S is therefore time in the same unit.

Example. A thermal fault is expected to cross a protected temperature in 40 min. Detection consumes 3 min, diagnosis 8 min, coordination 7 min and physical action 10 min. Slack = $40 - (3 + 8 + 7 + 10) = 12$ min.

Interpretation. Twelve minutes is not “spare time” for debate. It is the remaining margin for uncertainty, communication failure or a first action that does not work.

Limit. T_consequence is often uncertain. The safer practice is to calculate a range and use the lower credible value for urgent protective decisions.

Stale information is a governance hazard

Earth may send technically excellent advice that describes a configuration no longer present on Mars. Every delayed message should therefore carry the configuration, timestamp and assumptions on which it is based. A local decision record should do the same. The question is not only “Who said this?” but “For which state of the system was it true?”

Information age at decision

$$A_info = t_decision - t_measurement$$

If a pressure reading was measured at 14:05 and the decision occurs at 14:22, the information age is 17 minutes. Whether that is acceptable depends on how quickly the physical state can change. A 17-minute-old habitat pressure trend during a rapid leak may be nearly useless; a 17-minute-old inventory count may be adequate.

Operational rule. Every critical variable should have an acceptable information-age envelope tied to its dynamics.

Technical dissent needs a stopping rule

Dissent is valuable when it reveals an assumption, a sensor problem or a failure mode that the majority missed. It becomes paralysing when nobody knows how the discussion ends. The charter should therefore define a short structured challenge: state the disputed assumption, present the evidence, identify the consequence if the minority view is correct, and name the decision authority. For irreversible actions with large downside, a documented safety hold may be appropriate; for rapidly deteriorating conditions, the authority may have to act while preserving the dissent in the log for later review.

Handover is an engineering transfer

A command handover should transfer system state, protected limits, active waivers, degraded configurations, unresolved anomalies, next decision deadlines and current evidence. A verbal “you have command” is not enough. The incoming person should repeat back the critical state and confirm access to the tools and credentials needed to act.

Crisis governance drill

During a dust event, one engineer wants to stop a chemical process immediately because cooling margin is falling. The operations lead wants ten more minutes to finish a batch that produces a scarce reagent. Earth recommends continuing based on telemetry that is fifteen minutes old. Write the decision chain: what evidence is refreshed first, who may impose a safety stop, what reversible action buys time, and what is entered into the decision log?

► Reasoned solution

First-Man governance room: authority must be fast enough for physics and accountable enough for people

A Mars settlement cannot govern every decision by committee, nor can it survive by giving one commander unlimited discretion. Different decisions run on different clocks. A pressure leak may require action in seconds; a power-rationing plan may allow minutes; a change to food allocation or crew work rules may permit hours or days and should involve broader participation. Good operational governance assigns authority according to consequence and available decision time, then preserves traceability so emergency power does not quietly become permanent power.

The central design object is a decision-rights matrix. For each class of decision, define who may act immediately, who must be informed, what evidence must be recorded, when an independent review is required and what terminates emergency authority. This separates the speed needed for safety from the legitimacy needed for long-duration community life.

Decision slack shows when consultation is physically possible

$$S_{\text{dec}} = t_{\text{consequence}} - (t_{\text{detect}} + t_{\text{understand}} + t_{\text{act}})$$

Starting question

How much time remains for coordination or consultation after accounting for the minimum chain needed to detect, understand and act on a threat?

Read aloud

Read: “decision slack equals time to unacceptable consequence minus detection time, understanding time and action time.”

Symbols, pronunciation and meaning

$t_{\text{consequence}}$ is the time from event onset to unacceptable outcome without adequate control; t_{detect} is detection delay; $t_{\text{understand}}$ is time to classify the situation enough to choose an action; t_{act} is execution time; S_{dec} is remaining slack.

Units

All terms must use the same time unit, such as seconds or minutes.

Origin and status of values

These times come from hazard analysis, drills, observed operator performance and system response data. They are distributions in reality; a single number is a planning simplification.

Why this operation

The minimum response chain consumes part of the physical time available before consequence. The remainder is the time that can be spent on additional coordination without missing the safety window.

Substitution and calculation

Teaching case: unacceptable consequence in 12 min; detection 1 min; understanding 3 min; action 4 min. Slack = $12 - (1 + 3 + 4) = 4$ min.

Calculator entry

Enter $12 - (1 + 3 + 4)$.

Mental estimate

The response chain uses 8 of 12 minutes, leaving 4.

Independent check

$1 + 3 + 4 + 4 = 12$ min, reconstructing the consequence clock.

Physical or operational interpretation

Only four minutes remain for extra consultation in this simplified case. A governance process requiring a thirty-minute meeting is incompatible with the hazard.

Plain-English translation

Physics limits democracy at the second-to-second emergency timescale, but it does not justify bypassing accountability after the immediate danger has passed.

Variation / sensitivity

If detection improves from 1 min to 20 s, more slack becomes available. Better sensors and drills can therefore expand the space for deliberate decision-making.

Limit / assumption

The formula treats stages sequentially even though detection, interpretation and action can overlap. Use it as a conservative planning tool, not a universal law.

What this does not prove

Positive slack does not prove that a particular person should hold authority. Legitimacy, competence, conflict of interest and legal or mission rules remain separate questions.

Boundary case to test

If S_{dec} is negative, the current detection-and-response architecture is too slow for the hazard even before adding governance overhead. The solution is faster detection, simpler protective action or a safer system, not merely a more authoritarian decision rule.

Emergency authority needs an expiry mechanism

Every exceptional power should have a trigger, scope, recording requirement and return condition. “Until the crisis is over” is too vague. Better conditions are observable: atmosphere stable for a defined period, damaged section isolated, power reserve restored above a threshold, medical triage complete, or a formal review convened. This makes normal governance the default state rather than something that must be won back from emergency command.

Trigger

defined hazard and threshold

Act

bounded delegated authority

Record

evidence, dissent, orders, outcomes

Return

objective exit criteria and review

Technical dissent is a safety instrument

A team needs a protected way to say “stop” when a specialist believes the commander’s mental model is wrong. The dissent channel should specify who can request a pause, what happens if time is critical, how the concern is recorded and who performs the later review. This protects both safety and command: the commander can act quickly while the organisation preserves evidence that a warning existed and can learn from it.

Resource rationing needs a rule before scarcity

Water, power, EVA time, medical capability and communications bandwidth can all become scarce. A rationing framework should rank protection of life, prevention of irreversible system loss, restoration capability and then lower-consequence goals. The exact policy is a community and mission governance choice, but the ranking should be discussed in normal conditions. Inventing it during the emergency invites inconsistency and conflict.

Exercise: when should the commander stop deciding alone?

A dust event forces emergency power rationing. During the first ten minutes, automatic load shedding and the duty commander protect life support. After two hours the power state is stable, forecasts remain uncertain, and the question is whether to suspend a week of science and industrial activity. The immediate protective action belongs to emergency authority; the week-long allocation decision no longer runs on a seconds-scale hazard clock and should return to the settlement's broader governance process with technical evidence.

The closer a physical deadline becomes, the more governance must shift toward pre-authorised, evidence-based local action. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Crisis-governance qualification lab: make authority faster without making it arbitrary

A Mars settlement cannot wait for Earth to resolve every urgent decision, yet “local autonomy” is not a licence for undefined authority. Crisis governance should specify who may act, on which evidence, for how long, with what review and how normal authority is restored afterward. The objective is to reduce decision latency while preserving accountability and technical competence.

Pre-authorise decision domains

Before a crisis, define domains such as life-support protection, medical isolation, power shedding, evacuation, contamination quarantine and suspension of hazardous work. For each domain, record the normal decision owner, emergency substitute, consultation requirements, limits and review point. Some actions should be reversible by design; others, such as venting a compartment or consuming a strategic reserve, need a higher threshold because they cannot easily be undone.

Earth support remains valuable even with long communication delays. The settlement can transmit telemetry and decision logs, receive later analysis and update procedures. The key is to avoid a governance model in which urgent local safety decisions are frozen while waiting for external permission that cannot arrive within the hazard clock.

Decision time margin – full qualification formula

$$M_{\text{decision}} = t_{\text{hazard}} - (t_{\text{detect}} + t_{\text{understand}} + t_{\text{authorize}} + t_{\text{act}})$$

1 – Concrete question

Does the governance and response chain complete before the hazard crosses the unacceptable threshold?

2 – Intuition

The hazard has a clock. Detection, diagnosis, authorisation and action consume that clock.

3 – Quantities

Estimate the time to the critical threshold and the four major response delays.

4 – Formula

Decision margin equals hazard time minus total response time.

5 – Read aloud

"M decision equals t hazard minus t detect plus t understand plus t authorize plus t act."

6 – Symbols

Each t is a duration; M_{decision} is the remaining temporal margin.

7 – Pronunciation

Subscripts label stages in the decision chain.

8 – Units

Use minutes, hours or another single time unit.

9 – Convention

Positive margin indicates the planned chain fits within the assumed hazard clock; negative means the process is too slow.

10 – Why subtraction

Every response stage consumes time that is no longer available before the threshold.

11 – Assumptions

The hazard time estimate can be uncertain; conservative bounds are safer than optimistic averages.

12 – Unit check

$\text{min} - (\text{min} + \text{min} + \text{min} + \text{min}) = \text{min}.$

13 – Numerical case

Time until protected limit is crossed: $t_{\text{hazard}} = 45$ min.

Detection = 3 min; interpretation = 8 min; authorisation = 10 min; action = 9 min.

Decision/action chain = $3 + 8 + 10 + 9 = 30$ min.

$M_{\text{decision}} = 45 - 30 = 15$ min.

14 – Operations

Add response delays first. The seven-minute authorisation step is visible as a large fraction of the total.

15 – Algebra check

20 min response + 8 min margin = 28 min hazard clock.

16 – Mental estimate

The response consumes roughly three quarters of the clock, so less than ten minutes of margin is plausible.

17 – Interpretation

The chain works on paper, but governance latency is operationally significant.

18 – What it does not prove

It does not prove the chosen action is correct or that the hazard-time model is accurate.

19 – Sensitivity

If the authorised decision maker is unavailable and authorisation takes 12 minutes, margin falls to 3 minutes. A named deputy can therefore be a safety control.

20 – Practice

Guided exercise. Compute decision margin for a hazard in 45 minutes with stage times of 3, 8, 10 and 9 minutes.

Detailed guided correction.

1. Decision/action chain = $3 + 8 + 10 + 9 = 30$ min.
2. Margin = $45 - 30 = 15$ min.
3. The 15-minute margin belongs to this defined chain and hazard prediction; uncertainty in hazard progression should be represented separately.

Autonomous exercise. A reversible protective shutdown can be pre-authorised at the technical-controller level, reducing authorisation time from 10 to 2 minutes. Using the same 45-minute hazard, 3-minute detection, 8-minute interpretation and 9-minute action, compute the new margin and define one safeguard against misuse.

► Autonomous correction – open after attempting the exercise

21 – Mission decision

If governance latency consumes too much of the hazard clock, pre-authorise bounded protective actions and move later review after the immediate safety state is achieved.

Record dissent and uncertainty

Good crisis governance does not require artificial unanimity. The decision log should state what is known, what is uncertain, what alternatives were considered, who dissented and which observation would trigger reconsideration. This allows later review and reduces hindsight distortion. It also lets the next shift understand why a restriction exists instead of treating it as an unexplained order.

Qualification drill

Design authority for a water-contamination event that may require isolation of a habitat branch. Earth cannot answer before the first action is needed. Assign who can quarantine the branch, who can release it, what medical and engineering inputs are required, and when the decision must be reviewed. Then inject the temporary unavailability of the settlement commander and test whether the process still works.

Source context. NASA spaceflight operations and human-system standards provide relevant operational context. The governance model here is a Delta-Sierra educational construct, not NASA policy. [NASA JSC — Spaceflight Operations](#).

Resource restrictions need a rule for entry and exit

Emergency rationing becomes socially and operationally corrosive when nobody knows what evidence will end it. Every restriction should define why it was imposed, which metric is being protected, who can modify it and what measurable condition allows relaxation. This turns a political argument into an auditable operational state.

For example, a water restriction may begin because verified potable reserve falls below a defined number of days. The exit condition can require restored treatment capacity plus a target reserve, not simply “the repair is finished.” The same principle applies to power shedding, EVA suspension or quarantine.

Conflict between technical domains

Real crises produce legitimate competing priorities. Medical staff may want more power for a treatment area while the power team is protecting a depleted battery. Science may want rover access while maintenance needs the same vehicle. Governance should not pretend these conflicts disappear under command. It should provide a mechanism for exposing consequences in common terms—time to threshold, lives at risk, irreversibility, recovery cost—and naming the person authorised to choose when priorities remain incompatible.

Shift handover and decision continuity

A crisis can last longer than one team’s safe working period. Handover should preserve the evidence state, not merely the conclusion. The receiving shift needs observations, hypotheses, actions already taken, temporary configurations, decisions pending, dissenting views and trigger conditions. Without this record, each shift risks restarting diagnosis from zero or undoing a deliberate protective action.

Post-crisis review without hindsight theatre

After stabilisation, reconstruct the timeline using logs and telemetry before memories converge on a convenient story. Distinguish decisions that were reasonable with the information available at the time from outcomes that merely happened to be good or bad. The review should produce specific changes to procedures, training, instrumentation or authority boundaries. Blame without mechanism does not improve the settlement.

R59 authority-under-delay board: make protective action fast, bounded and reviewable

Governance in a Mars settlement is not only constitutional design. During a fast technical event it becomes a timing problem: who may act, on what evidence, within what limits, and how is the decision handed over or reversed? Communication delay means that some decisions must be made locally even when Earth expertise is valuable. The answer should be neither unrestricted command authority nor a chain so slow that the hazard wins.

Pre-authorise reversible protection

Define actions that can be taken rapidly because they are bounded and reversible: shedding a discretionary load, isolating a suspect branch, pausing an EVA launch, switching to a protected reserve or moving people into a refuge. For each action, specify entry conditions, authority, automatic logging, maximum duration and who must review continuation. This converts “act fast” into a controlled operating rule.

Reserve higher authority for irreversible or rights-limiting decisions

Permanent equipment abandonment, destructive isolation, severe resource restrictions or disciplinary measures can have long consequences. Their authority path should be explicit and should include technical evidence, affected-domain input and a record of uncertainty.

Emergency rules may shorten the path, but they should also define when emergency authority expires.

Dissent is an operational sensor

A technically qualified person who disagrees with the dominant interpretation may be detecting a weak signal. Recording dissent does not mean paralysing the decision. The incident log can capture the alternative hypothesis, evidence that would support it and the condition that would trigger reconsideration. This makes later review more honest and can protect against groupthink during fatigue.

Handover must transfer unresolved questions

Shift turnover should not be a list of completed actions. It should carry current configuration, temporary limits, active hypotheses, pending tests, authority already exercised, next decision deadline and the explicit “if/then” triggers that could change the plan. A fresh team that sees only the latest stable telemetry can otherwise repeat a discarded diagnosis or remove a protection whose purpose is no longer obvious.

Decision drill — use the extra eight minutes

The autonomous exercise shows how pre-authorising a reversible shutdown can increase decision margin from 15 to 23 minutes. Spend that margin deliberately: obtain a confirming measurement, prepare the recovery path, communicate with affected users or preserve a refuge option. A faster authority path is valuable only if the saved time improves safety rather than encouraging delay until the last possible moment.

Primary-source bridge. NASA’s spaceflight-operations organisation and Human Research Program provide context for operational decision-making and human performance. The authority framework here is a Delta-Sierra educational design, not a NASA governance rule. [NASA JSC — Spaceflight Operations](#).

R60 crisis governance: authority must move faster than the hazard without becoming arbitrary

Operational governance on Mars is not mainly about writing a constitution. During a fast-moving technical event, governance is the mechanism that converts evidence into an authorised action before a protected limit is crossed. A good design therefore names who can isolate a system, evacuate a zone, consume emergency reserves, cancel an EVA, quarantine inventory or override a schedule—and also defines when that authority transfers if the nominal decision-maker is unavailable.

Pre-authorise bounded emergency actions

Waiting for a committee can be more dangerous than a technically imperfect but timely response. Teams should pre-authorise narrow actions when objective triggers are reached: close a valve, trip a feeder, stop a chemical process, shelter for radiation, recall an EVA crew or

isolate a habitat compartment. The authority should be bounded by the trigger and documented afterward. This reduces delay without giving one person unlimited discretionary power.

Primary-source bridge. NASA's spaceflight operations material provides context for disciplined mission operations, roles and procedures. The settlement governance model here extends that operational logic to a long-duration surface community. [NASA — Spaceflight Operations](#).

Separate technical recommendation from command authority

The best technical expert may not be the person who owns the system-wide consequence. During a power crisis, an electrical specialist can recommend shedding a branch, while the duty commander decides whether the loss of that branch is acceptable for medical, thermal or life-support functions. This separation should not create bureaucracy; it should ensure that local optimisation does not accidentally damage a higher-priority function.

Use decision clocks that are tied to physical limits

Every urgent event should have a decision clock based on the earliest protected limit: time to oxygen threshold, battery depletion, toxic exposure, overheating, radiation shelter deadline or rescue return reserve. The clock turns “urgent” into a measurable constraint. If the remaining time is shorter than the normal approval chain, authority must already have a degraded-mode path. Otherwise the organisation has designed a procedure that cannot complete before the hazard wins.

Primary source at use. NASA JSC Spaceflight Operations is the operational bridge for time-bounded procedures, mission-control discipline and crew/ground decision interfaces. The specific Mars decision-clock model remains a Delta-Sierra teaching construction. [NASA JSC Spaceflight Operations](#).

Evidence thresholds should change with reversibility

A reversible low-cost action can be taken on weaker evidence than an irreversible action that destroys scarce hardware or commits the settlement to a new risk. Governance should therefore ask two questions: how certain are we, and how reversible is the proposed action? Closing an

isolation valve can be immediately reversible; venting a resource, abandoning a vehicle or consuming a unique spare may not be. The evidence threshold should rise with irreversibility and consequence.

Degraded communications require local authority

Earth advice may be delayed, unavailable or outdated relative to the local event. Each critical system should therefore have local decision rights and locally accessible procedures. Remote expertise is valuable for diagnosis and recovery planning, but it cannot be a mandatory dependency for time-critical protection. The settlement should practise operating with delayed external support so that the first communications outage is not also the first time authority is exercised locally.

Record the decision trail without slowing the response

During the acute phase, logging should be lightweight: time, observation, action, authority and immediate reason. After stabilisation, expand the record with data snapshots, configuration state, alternative options considered and consequences. This produces evidence for learning and accountability without forcing operators to write a report while a protected limit is approaching.

Scenario exercise — correct expertise, wrong authority chain

A chemical sensor trend indicates a hazardous release could cross a protected limit in 18 minutes. The technical specialist needs 4 minutes to validate the sensor and 3 minutes to recommend isolation. The normal approval chain requires two sequential reviews averaging 8 minutes each. That chain cannot finish inside the physical decision window. The governance defect exists before the emergency: the settlement needs a pre-authorised isolation rule or a shorter degraded-mode authority path. Training should score whether the team recognises the governance failure, not merely whether it calculates $4 + 3 + 16 = 23$ minutes.

Primary-source bridge. NASA-STD-3001 Volume 2 addresses human-system considerations that remain relevant to crew performance and safety. The authority structure in this lesson is a Delta-Sierra operational teaching model, not a NASA governance prescription. [NASA-STD-3001 Volume 2](#).

R60 governance drill: conflict between two locally correct decisions

Present the settlement with a power shortage during which the electrical team recommends shedding a district immediately, while the medical team warns that the same district contains a patient whose treatment depends on powered equipment. Both local recommendations can be technically correct. The governance task is to identify the protected functions, time-to-limit for each, alternatives, authority and the least irreversible action that preserves both where possible.

The exercise should score communication quality as well as the final decision. Did each team state observations separately from assumptions? Did they identify what information would change the decision? Was a decision deadline explicit? Did the commander know which action was reversible and which consumed an emergency reserve? These questions make governance observable and trainable rather than philosophical.

After the event, conduct a short decision reconstruction. Compare the timeline operators believed they had with the physical timeline later derived from data. If the organisation systematically underestimates approval or communication time, the fix may be procedural rather than technical. Governance improves when near-misses become measured evidence about the organisation itself.

Crisis governance is incomplete until temporary authority is reconciled, decisions are recorded and normal accountability resumes. Pedagogical synthesis by Delta-Sierra from the primary sources cited in this course; schematic, not to scale.

Primary sources and bridges

- [NASA JSC — Spaceflight Operations](#)
- [NASA — Moon to Mars Architecture](#)
- [NASA-STD-3001 Volume 2](#)