

Maintenance, fleet reliability and repair on Mars

Design a base that keeps operating as components age, fail and must be diagnosed, cannibalized, repaired or replaced without an immediate supply chain.

Before you begin — Prerequisites: modules 00–34 as relevant. Every important symbol is defined at first use.

Mastery objectives

- connect principles to architecture or operational decisions
- repeat simple calculations and verify units and assumptions
- identify degraded modes, interfaces and uncertainty
- produce a verifiable procedure or plan

1. Reliability and maintainability are not the same thing

Reliability describes the probability that equipment performs its function for a given period. Maintainability describes how easily and quickly it can be diagnosed, opened, repaired, tested and returned to service. A highly reliable but irreparable system can become dangerous on Mars; a somewhat less reliable but modular and repairable system may deliver better operational availability.

2. Design for access, not performance alone

A filter hidden behind three assemblies can turn a ten-minute task into a six-hour intervention. Connectors, access panels, handling masses, torque requirements, tools and work zones must be designed from the start. Common interfaces allow one assembly to replace another without rebuilding the whole system.

3. From symptom to cause: troubleshooting and fault trees

Troubleshooting separates symptom, failure mechanism and root cause. A falling water flow can be caused by a pump, a clogged filter, a partly closed valve, a bad sensor or a leak. Crew need telemetry, measurement points and procedures that prevent healthy hardware from being replaced by guesswork.

4. Repair at the lowest reasonable level

On Earth an entire module may be replaced because logistics make that convenient. On Mars spare mass is limited. Repairing a board, motor, seal or pump can save tens of kilograms of inventory. Going too deep, however, increases training, tools and error risk. The correct repair level is an architecture trade.

5. Redundancy, cannibalization and technical debt

Identical equipment provides redundancy and can enable cannibalization: one unavailable asset becomes a donor for another. The strategy must be controlled because it creates technical debt. Every removed part must be recorded so the base knows which assets are no longer complete.

6. Preventive, predictive and corrective maintenance

Preventive maintenance is scheduled by time or operating hours. Predictive maintenance uses vibration, temperature, current or pressure trends to intervene when degradation begins. Corrective maintenance occurs after failure. Mars operations must blend all three to reduce both spare mass and sudden-failure risk.

7. Common-cause failures and genuine independence

Two identical units do not automatically create robust redundancy. If they share the same power converter, software defect, environmental sensor or maintenance error, a single cause can disable both. Reliability work must therefore distinguish independent failures from common-cause failures. On a Mars base, two circulation pumps mounted together might both be lost if fluid contamination, local overheating or an incorrect maintenance procedure affects the whole subsystem. Useful redundancy may require separated power paths, diverse sensing, isolation capability, degraded modes and procedures that confirm the backup channel is actually healthy before load is transferred to it.

8. Sizing spares from risk rather than fear

Carrying a duplicate of every part sounds reassuring but quickly becomes impossible as mission duration, mass and volume grow. Spare provisioning must combine failure frequency, criticality, replacement time, local repair capability, parts commonality and the consequences of stockout. A light but unique seal may justify several copies, while a heavy and highly reliable structure may be better served by a repair strategy than by full replacement. Maintainers must also consider hidden consumables such as lubricants, filters, connectors, fasteners, adhesives, cleaning materials and test equipment. Maintenance logistics therefore becomes a quantitative extension of reliability engineering rather than a list written after design is complete.

9. Closing the loop: failure, lessons learned and configuration change

A base that repairs equipment without learning will accumulate recurring failures. Every significant incident should create an operational record: initial symptoms, telemetry, hypotheses tested, accepted cause, parts replaced, post-repair condition and any procedural change. If the same failure returns, the team must decide whether the root problem is maintenance, environment or design. Permanent changes then have to pass through configuration management so drawings, software, inventories and procedures continue to describe the hardware that actually exists on Mars. Maintenance therefore becomes a continuous-improvement process, not merely a way to restore service.

10. Worked example: availability of a critical function

Assume equipment operates an average 2,000 h between failures and requires 20 h to repair. Approximate availability is $A = \text{MTBF}/(\text{MTBF} + \text{MTTR}) = 2,000/(2,000 + 20) \approx 0.9901$, or 99.01%. If repair time rises to 200 h because access or parts are poor, availability drops to 90.9%. Maintainability can therefore cost almost ten percentage points without changing intrinsic reliability.

Deeper engineering: size availability rather than counting failures alone

A life-critical system can fail rarely and still be unavailable for too long if repair is slow. Availability therefore connects two different quantities: mean time between failures and mean time to restore the function. NASA treats reliability and maintainability as related but distinct disciplines. On Mars the distinction becomes operational because the workshop, technician or spare part may itself be unavailable. [NASA — Reliability and Maintainability](#)

Worked example. Assume a compressor with a mean time between failures of 2,000 hours and a mean repair time of 20 hours. Simplified intrinsic availability is $A = \text{MTBF} \div (\text{MTBF} + \text{MTTR})$. Here $A = 2,000 \div (2,000 + 20) = 2,000 \div 2,020 \approx 0.9901$, or about 99.01%. That sounds excellent, but over an 8,760-hour year, 0.99% unavailability is still about 86.7 hours.

A life-critical function must therefore be judged together with redundancy and degraded modes, not by the percentage alone.

The same equation shows where engineering effort can help. If compressor reliability is unchanged but access, procedures and spares reduce MTTR from 20 hours to 5 hours, $A = 2,000 \div 2,005 \approx 99.75\%$. The improvement comes from maintainability. A settlement must therefore invest in access, diagnosis, standardisation, tooling and training as deliberately as it invests in robust hardware. [NASA NTRS — Supportability Concepts](#)

11. Progressive exercise

A base owns three identical pumps. Two are required for nominal operation and one is a spare. Propose a rotation, vibration-monitoring and common-spares strategy. Explain what changes if the main seal has a 30% probability of failing during a 500-day interval.

Reasoned correction

One defensible strategy periodically rotates pumps A, B and C so the spare does not remain untouched for hundreds of days: two operate while the third is inspected, then service is redistributed using operating hours and condition. Vibration trends should be recorded with flow, temperature and motor current to separate mechanical wear from load changes. If each main seal independently has a 30% chance of failure over 500 days, expected failures across three pumps are 0.9 and the probability of at least one failure is $1 - 0.7^3 \approx 65.7\%$. That supports carrying seals, replacement tooling and enhanced monitoring; a common-cause mechanism would make the independence assumption even more optimistic.

Mini-project

Build a supportability plan for a Mars workshop: critical functions, repair levels, tools, offline documentation, common parts, cannibalization criteria, return-to-service tests and availability indicators.

Maintenance is the system that keeps every other system alive

A Mars settlement cannot treat maintenance as a support activity that begins after hardware breaks. Maintenance is part of the architecture because every pump, valve, rover, fan, suit component and sensor has a finite life and a repair burden. The design question is therefore not only whether a device works, but whether the crew can detect degradation, isolate it, reach the failed item, remove it, install a replacement, test the repair and return the system to service with the tools and documentation available locally.

Fleet thinking is especially important. Ten identical pumps can simplify training and spares, yet a common design defect can threaten all ten. A mixed fleet may reduce common-cause risk but increases tools, documentation and inventory. The settlement should decide where standardization helps and where diversity is worth its logistics penalty. Those decisions are better made from failure consequences and repair data than from a blanket preference for one philosophy.

Good maintenance also protects knowledge. A replacement performed without recording the symptom, removed part, configuration and test result solves only today's problem. A recorded repair becomes reliability evidence. Over hundreds of interventions, that evidence reveals wear patterns, bad lots, weak procedures and components that deserve redesign. On Mars, the maintenance log is therefore part of the engineering database.

Four reliability ideas a repair team must use correctly

MTBF

Mean Time Between Failures is an average interval between failures for repairable equipment under a stated population and operating context. It is not a promise that a particular unit will run for exactly that many hours.

MTTR

Mean Time To Repair is an average restoration time under declared maintenance conditions. It can hide large variation caused by access, diagnosis, missing tools or parts, so crews should also study the repair-time distribution and worst credible cases.

Availability

Availability is the fraction of time a function is capable of being used when required. It combines how often failures occur with how long restoration takes, so improving repair time can sometimes be as valuable as improving component life.

Cannibalization

Cannibalization removes a usable component from one item to restore another. It can be rational during a shortage but creates configuration, traceability and future fleet-capacity consequences that must be controlled.

Calculation laboratory

Formula 1 — simple inherent availability estimate

Canonical quantitative mini-lessons

This block is the module quantitative reference; exercise values are pedagogical and must be replaced by mission data before any real decision.

Intrinsic availability

$$A = \text{MTBF} / (\text{MTBF} + \text{MTTR})$$

1 — Concrete question

What does “ $A = \text{MTBF} / (\text{MTBF} + \text{MTTR})$ ” compute in “Intrinsic availability”?

2 — Intuition without symbols

Availability depends on both failure frequency and restoration speed.

3 — Quantities

A: intrinsic availability [sans dimension]; MTBF: mean time between failures [h]; MTTR: mean time to repair [h]

4 — Formula

$$A = \text{MTBF} / (\text{MTBF} + \text{MTTR})$$

5 — Read aloud

Read “ $A = \text{MTBF} / (\text{MTBF} + \text{MTTR})$ ” by naming every operation, subscript and grouping explicitly.

6 — Symbols and meaning

A: intrinsic availability [sans dimension]; MTBF: mean time between failures [h]; MTTR: mean time to repair [h]

7 — Pronunciation

Say subscripts, powers, parentheses and Greek letters explicitly so the relation can be reconstructed without seeing the page.

8 — Units

A [sans dimension]; MTBF [h]; MTTR [h]

9 — Convention

Keep the same frame, time basis, system boundary and sign convention throughout the substitution.

10 — Why this operation

Division relates a quantity to a coherent capacity, duration or reference.

11 — Assumptions

Numerator and denominator must describe the same scenario and the denominator must remain positive.

12 — Independent check

Multiplying the result by the denominator should reconstruct the numerator.

13 — Numerical case

With MTBF = 1000 h, MTTR = 10 h: $A = 1000 / (1000 + 10) = 0.9901$.

14 — Why the calculation works

Division relates a quantity to a coherent capacity, duration or reference.

15 — Verification

Multiplying the result by the denominator should reconstruct the numerator.

16 — Mental estimate

Comparing numerator and denominator scales predicts the result.

17 — Interpretation

Improving maintainability can matter as much as increasing MTBF when downtime is critical.

18 — What the result does not prove

The calculation proves neither that inputs are exact nor that the model remains valid outside the declared scenario; both require separate evidence.

19 — Sensitivity

The result rises with the numerator and falls as the denominator rises.

20 — Guided and autonomous exercises

Guided exercise. Recalculate this scenario: With MTBF = 500 h, MTTR = 20 h: $A = 500 / (500 + 20)$?

Detailed guided correction — open after trying

With MTBF = 500 h, MTTR = 20 h: $A = 500 / (500 + 20) = 0.9615$. The decision must then be checked against the module margins and assumptions.

Autonomous exercise. Recalculate this scenario: With MTBF = 2000 h, MTTR = 5 h: $A = 2000 / (2000 + 5)$?

Autonomous correction — open after trying

With MTBF = 2000 h, MTTR = 5 h: $A = 2000 / (2000 + 5) = 0.9975$. The decision must then be checked against the module margins and assumptions.

21 — Mission decision

Improving maintainability can matter as much as increasing MTBF when downtime is critical.

Expected failures

$N_{\text{fail}} = \text{fleet_hours} / \text{MTBF}$

1 — Concrete question

What does “ $N_{\text{fail}} = \text{fleet_hours} / \text{MTBF}$ ” compute in “Expected failures”?

2 — Intuition without symbols

Cumulative fleet hours divided by MTBF give an order-of-magnitude count of events to handle.

3 — Quantities

N_{fail} : expected failures [failure]; fleet_hours: cumulative operating hours [h]; MTBF: mean time between failures [h/failure]

4 — Formula

$N_{\text{fail}} = \text{fleet_hours} / \text{MTBF}$

5 — Read aloud

Read “ $N_{\text{fail}} = \text{fleet_hours} / \text{MTBF}$ ” by naming every operation, subscript and grouping explicitly.

6 — Symbols and meaning

N_{fail} : expected failures [failure]; fleet_hours: cumulative operating hours [h]; MTBF: mean time between failures [h/failure]

7 — Pronunciation

Say subscripts, powers, parentheses and Greek letters explicitly so the relation can be reconstructed without seeing the page.

8 — Units

N_{fail} [failure]; fleet_hours [h]; MTBF [h/failure]

9 — Convention

Keep the same frame, time basis, system boundary and sign convention throughout the substitution.

10 — Why this operation

Division relates a quantity to a coherent capacity, duration or reference.

11 — Assumptions

Numerator and denominator must describe the same scenario and the denominator must remain positive.

12 — Independent check

Multiplying the result by the denominator should reconstruct the numerator.

13 — Numerical case

With fleet_hours = 5000 h, MTBF = 1000 h/failure: $N_{\text{fail}} = 5000 / 1000 = 5$ failure.

14 — Why the calculation works

Division relates a quantity to a coherent capacity, duration or reference.

15 — Verification

Multiplying the result by the denominator should reconstruct the numerator.

16 — Mental estimate

Comparing numerator and denominator scales predicts the result.

17 — Interpretation

Size crew and spares for total fleet exposure rather than a single unit.

18 — What the result does not prove

The calculation proves neither that inputs are exact nor that the model remains valid outside the declared scenario; both require separate evidence.

19 — Sensitivity

The result rises with the numerator and falls as the denominator rises.

20 — Guided and autonomous exercises

Guided exercise. Recalculate this scenario: With fleet_hours = 2400 h, MTBF = 800 h/failure: $N_{\text{fail}} = 2400 / 800$?

Detailed guided correction — open after trying

With fleet_hours = 2400 h, MTBF = 800 h/failure: $N_{\text{fail}} = 2400 / 800 = 3$ failure. The decision must then be checked against the module margins and assumptions.

Autonomous exercise. Recalculate this scenario: With fleet_hours = 10000 h, MTBF = 2000 h/failure: $N_{\text{fail}} = 10000 / 2000$?

Autonomous correction — open after trying

With fleet_hours = 10000 h, MTBF = 2000 h/failure: $N_{\text{fail}} = 10000 / 2000 = 5$ failure. The decision must then be checked against the module margins and assumptions.

21 — Mission decision

Size crew and spares for total fleet exposure rather than a single unit.

Maintenance workload

$$H_{\text{maint}} = N_{\text{fail}} \times \text{MTTR}$$

1 — Concrete question

What does " $H_{\text{maint}} = N_{\text{fail}} \times \text{MTTR}$ " compute in "Maintenance workload"?

2 — Intuition without symbols

Each failure consumes repair time; their product gives a first maintenance workload.

3 — Quantities

H_{maint} : maintenance hours [h]; N_{fail} : expected failures [failure]; MTTR: mean repair time [h/failure]

4 — Formula

$$H_{\text{maint}} = N_{\text{fail}} \times \text{MTTR}$$

5 — Read aloud

Read " $H_{\text{maint}} = N_{\text{fail}} \times \text{MTTR}$ " by naming every operation, subscript and grouping explicitly.

6 — Symbols and meaning

H_{maint} : maintenance hours [h]; N_{fail} : expected failures [failure]; MTTR: mean repair time [h/failure]

7 — Pronunciation

Say subscripts, powers, parentheses and Greek letters explicitly so the relation can be reconstructed without seeing the page.

8 — Units

H_{maint} [h]; N_{fail} [failure]; MTTR [h/failure]

9 — Convention

Keep the same frame, time basis, system boundary and sign convention throughout the substitution.

10 — Why this operation

Multiplication converts a rate, fraction or specific quantity into a total.

11 — Assumptions

Factors must belong to the same scenario and their units must combine correctly.

12 — Independent check

Dividing the result by a non-zero factor should recover the product of the others.

13 — Numerical case

With $N_{\text{fail}} = 5$ failure, $\text{MTTR} = 8$ h/failure: $H_{\text{maint}} = 5 \times 8 = 40$ h.

14 — Why the calculation works

Multiplication converts a rate, fraction or specific quantity into a total.

15 — Verification

Dividing the result by a non-zero factor should recover the product of the others.

16 — Mental estimate

Estimating each factor separately quickly predicts product scale.

17 — Interpretation

Compare this workload with technician time actually available before accepting the fleet plan.

18 — What the result does not prove

The calculation proves neither that inputs are exact nor that the model remains valid outside the declared scenario; both require separate evidence.

19 — Sensitivity

A relative change in one factor transfers directly to the product when the others stay fixed.

20 — Guided and autonomous exercises

Guided exercise. Recalculate this scenario: With $N_{\text{fail}} = 3$ failure, $\text{MTTR} = 12$ h/failure: $H_{\text{maint}} = 3 \times 12$?

Detailed guided correction — open after trying

With $N_{\text{fail}} = 3$ failure, $\text{MTTR} = 12$ h/failure: $H_{\text{maint}} = 3 \times 12 = 36$ h. The decision must then be checked against the module margins and assumptions.

Autonomous exercise. Recalculate this scenario: With $N_{\text{fail}} = 10$ failure, $\text{MTTR} = 4$ h/failure: $H_{\text{maint}} = 10 \times 4$?

Autonomous correction — open after trying

With $N_{\text{fail}} = 10$ failure, $\text{MTTR} = 4$ h/failure: $H_{\text{maint}} = 10 \times 4 = 40$ h. The decision must then be checked against the module margins and assumptions.

21 — Mission decision

Compare this workload with technician time actually available before accepting the fleet plan.

Spare-stock coverage

$$t_{\text{spares}} = N_{\text{spares}} / \lambda_{\text{fail}}$$

1 — Concrete question

What does " $t_{\text{spares}} = N_{\text{spares}} / \lambda_{\text{fail}}$ " compute in "Spare-stock coverage"?

2 — Intuition without symbols

A spare stock becomes a duration when compared with average consumption rate.

3 — Quantities

t_{spares} : coverage duration [d]; N_{spares} : available spares [spare]; λ_{fail} : average spare consumption [spare/d]

4 — Formula

$t_{\text{spares}} = N_{\text{spares}} / \lambda_{\text{fail}}$

5 — Read aloud

Read " $t_{\text{spares}} = N_{\text{spares}} / \lambda_{\text{fail}}$ " by naming every operation, subscript and grouping explicitly.

6 — Symbols and meaning

t_{spares} : coverage duration [d]; N_{spares} : available spares [spare]; λ_{fail} : average spare consumption [spare/d]

7 — Pronunciation

Say subscripts, powers, parentheses and Greek letters explicitly so the relation can be reconstructed without seeing the page.

8 — Units

t_{spares} [d]; N_{spares} [spare]; λ_{fail} [spare/d]

9 — Convention

Keep the same frame, time basis, system boundary and sign convention throughout the substitution.

10 — Why this operation

Division relates a quantity to a coherent capacity, duration or reference.

11 — Assumptions

Numerator and denominator must describe the same scenario and the denominator must remain positive.

12 — Independent check

Multiplying the result by the denominator should reconstruct the numerator.

13 — Numerical case

With $N_{\text{spares}} = 12$ spare, $\lambda_{\text{fail}} = 0.5$ spare/d: $t_{\text{spares}} = 12 / 0.5 = 24$ day.

14 — Why the calculation works

Division relates a quantity to a coherent capacity, duration or reference.

15 — Verification

Multiplying the result by the denominator should reconstruct the numerator.

16 — Mental estimate

Comparing numerator and denominator scales predicts the result.

17 — Interpretation

Resupply or manufacture locally before coverage falls below realistic logistics delay.

18 — What the result does not prove

The calculation proves neither that inputs are exact nor that the model remains valid outside the declared scenario; both require separate evidence.

19 — Sensitivity

The result rises with the numerator and falls as the denominator rises.

20 — Guided and autonomous exercises

Guided exercise. Recalculate this scenario: With $N_{\text{spares}} = 8$ spare, $\lambda_{\text{fail}} = 0.25$ spare/d: $t_{\text{spares}} = 8 / 0.25$?

Detailed guided correction — open after trying

With $N_{\text{spares}} = 8$ spare, $\lambda_{\text{fail}} = 0.25$ spare/d: $t_{\text{spares}} = 8 / 0.25 = 32$ day. The decision must then be checked against the module margins and assumptions.

Autonomous exercise. Recalculate this scenario: With $N_{\text{spares}} = 20$ spare, $\lambda_{\text{fail}} = 1$ spare/d: $t_{\text{spares}} = 20 / 1$?

Autonomous correction — open after trying

With $N_{\text{spares}} = 20$ spare, $\lambda_{\text{fail}} = 1$ spare/d: $t_{\text{spares}} = 20 / 1 = 20$ day. The decision must then be checked against the module margins and assumptions.

21 — Mission decision

Resupply or manufacture locally before coverage falls below realistic logistics delay.

Fleet availability

$A_{\text{fleet}} = n_{\text{ready}} / n_{\text{total}}$

1 — Concrete question

What does " $A_{\text{fleet}} = n_{\text{ready}} / n_{\text{total}}$ " compute in "Fleet availability"?

2 — Intuition without symbols

Available fleet is the fleet that can actually be committed now.

3 — Quantities

A_{fleet} : ready fraction [sans dimension]; n_{ready} : ready units [unit]; n_{total} : total units [unit]

4 — Formula

$A_{\text{fleet}} = n_{\text{ready}} / n_{\text{total}}$

5 — Read aloud

Read " $A_{\text{fleet}} = n_{\text{ready}} / n_{\text{total}}$ " by naming every operation, subscript and grouping explicitly.

6 — Symbols and meaning

A_{fleet} : ready fraction [sans dimension]; n_{ready} : ready units [unit]; n_{total} : total units [unit]

7 — Pronunciation

Say subscripts, powers, parentheses and Greek letters explicitly so the relation can be reconstructed without seeing the page.

8 — Units

A_{fleet} [sans dimension]; n_{ready} [unit]; n_{total} [unit]

9 — Convention

Keep the same frame, time basis, system boundary and sign convention throughout the substitution.

10 — Why this operation

Division relates a quantity to a coherent capacity, duration or reference.

11 — Assumptions

Numerator and denominator must describe the same scenario and the denominator must remain positive.

12 — Independent check

Multiplying the result by the denominator should reconstruct the numerator.

13 — Numerical case

With $n_{\text{ready}} = 9$ unit, $n_{\text{total}} = 10$ unit: $A_{\text{fleet}} = 9 / 10 = 0.9$.

14 — Why the calculation works

Division relates a quantity to a coherent capacity, duration or reference.

15 — Verification

Multiplying the result by the denominator should reconstruct the numerator.

16 — Mental estimate

Comparing numerator and denominator scales predicts the result.

17 — Interpretation

Suspend non-critical tasks when ready fraction threatens life-critical redundancy.

18 — What the result does not prove

The calculation proves neither that inputs are exact nor that the model remains valid outside the declared scenario; both require separate evidence.

19 — Sensitivity

The result rises with the numerator and falls as the denominator rises.

20 — Guided and autonomous exercises

Guided exercise. Recalculate this scenario: With $n_{\text{ready}} = 14$ unit, $n_{\text{total}} = 16$ unit: $A_{\text{fleet}} = 14 / 16$?

Detailed guided correction — open after trying

With $n_{\text{ready}} = 14$ unit, $n_{\text{total}} = 16$ unit: $A_{\text{fleet}} = 14 / 16 = 0.875$. The decision must then be checked against the module margins and assumptions.

Autonomous exercise. Recalculate this scenario: With $n_{\text{ready}} = 6$ unit, $n_{\text{total}} = 8$ unit: $A_{\text{fleet}} = 6 / 8$?

Autonomous correction — open after trying

With $n_{\text{ready}} = 6$ unit, $n_{\text{total}} = 8$ unit: $A_{\text{fleet}} = 6 / 8 = 0.75$. The decision must then be checked against the module margins and assumptions.

21 — Mission decision

Suspend non-critical tasks when ready fraction threatens life-critical redundancy.

Series-chain reliability

$$R_{\text{series}} = R_1 \times R_2 \times R_3$$

1 — Concrete question

What does “ $R_{\text{series}} = R_1 \times R_2 \times R_3$ ” compute in “Series-chain reliability”?

2 — Intuition without symbols

In a series chain every function must succeed; reliabilities multiply.

3 — Quantities

R_{series} : chain reliability [sans dimension]; R_1 : component one reliability [sans dimension]; R_2 : component two reliability [sans dimension]; R_3 : component three reliability [sans dimension]

4 — Formula

$$R_{\text{series}} = R_1 \times R_2 \times R_3$$

5 — Read aloud

Read “ $R_{\text{series}} = R_1 \times R_2 \times R_3$ ” by naming every operation, subscript and grouping explicitly.

6 — Symbols and meaning

R_{series} : chain reliability [sans dimension]; R_1 : component one reliability [sans dimension]; R_2 : component two reliability [sans dimension]; R_3 : component three reliability [sans dimension]

7 — Pronunciation

Say subscripts, powers, parentheses and Greek letters explicitly so the relation can be reconstructed without seeing the page.

8 — Units

R_{series} [sans dimension]; R_1 [sans dimension]; R_2 [sans dimension]; R_3 [sans dimension]

9 — Convention

Keep the same frame, time basis, system boundary and sign convention throughout the substitution.

10 — Why this operation

Multiplication converts a rate, fraction or specific quantity into a total.

11 — Assumptions

Factors must belong to the same scenario and their units must combine correctly.

12 — Independent check

Dividing the result by a non-zero factor should recover the product of the others.

13 — Numerical case

With $R_1 = 0.99$ sans dimension, $R_2 = 0.98$ sans dimension, $R_3 = 0.97$ sans dimension: $R_{\text{series}} = 0.99 \times 0.98 \times 0.97 = 0.9411$.

14 — Why the calculation works

Multiplication converts a rate, fraction or specific quantity into a total.

15 — Verification

Dividing the result by a non-zero factor should recover the product of the others.

16 — Mental estimate

Estimating each factor separately quickly predicts product scale.

17 — Interpretation

Identify the link dominating the loss instead of hiding weakness behind the others.

18 — What the result does not prove

The calculation proves neither that inputs are exact nor that the model remains valid outside the declared scenario; both require separate evidence.

19 — Sensitivity

A relative change in one factor transfers directly to the product when the others stay fixed.

20 — Guided and autonomous exercises

Guided exercise. Recalculate this scenario: With $R_1 = 0.95$ sans dimension, $R_2 = 0.95$ sans dimension, $R_3 = 0.95$ sans dimension: $R_{\text{series}} = 0.95 \times 0.95 \times 0.95$?

Detailed guided correction — open after trying

With $R_1 = 0.95$ sans dimension, $R_2 = 0.95$ sans dimension, $R_3 = 0.95$ sans dimension: $R_{\text{series}} = 0.95 \times 0.95 \times 0.95 = 0.8574$. The decision must then be checked against the module margins and assumptions.

Autonomous exercise. Recalculate this scenario: With $R_1 = 0.999$ sans dimension, $R_2 = 0.99$ sans dimension, $R_3 = 0.98$ sans dimension: $R_{\text{series}} = 0.999 \times 0.99 \times 0.98$?

Autonomous correction — open after trying

With $R_1 = 0.999$ sans dimension, $R_2 = 0.99$ sans dimension, $R_3 = 0.98$ sans dimension: $R_{\text{series}} = 0.999 \times 0.99 \times 0.98 = 0.9692$. The decision must then be checked against the module margins and assumptions.

21 — Mission decision

Identify the link dominating the loss instead of hiding weakness behind the others.

Starting question

What fraction of time would a repairable item be available in a simplified failure-and-repair model?

Read aloud

Say: “availability equals mean time between failures divided by mean time between failures plus mean time to repair.”

Symbols, pronunciation and meaning

A is availability, MTBF is mean operating time between failures and MTTR is mean restoration time.

Units

MTBF and MTTR must use the same time unit. Their ratio is dimensionless and is often reported as a percentage.

Origin and status of values

Values should come from relevant fleet data or clearly labelled planning assumptions under comparable duty and maintenance conditions.

Why this operation

The denominator represents one simplified operating-plus-repair cycle; the numerator is the portion of that cycle spent available.

Substitution and calculation

With MTBF = 1,000 h and MTTR = 20 h: $A = 1,000 / 1,020 = 0.9804$, or about 98.0%.

Calculator entry

Enter $1000 \div (1000 + 20)$, then multiply by 100 only for percentage display.

Mental estimate

Repair occupies 20 of roughly 1,020 hours, about 2%, so availability should be near 98%.

Independent check

Unavailability is about $1 - 0.9804 = 0.0196$. Multiplying that fraction by 1,020 h gives about 20 h.

Physical or operational interpretation

A high availability figure can still be inadequate if the unavailable periods occur during a critical life-support demand with no redundancy.

Plain-English translation

In this simplified cycle, the item is capable of service about ninety-eight percent of the time.

Variation / sensitivity

Reducing MTTR from 20 h to 5 h raises availability to about 99.5% without changing the failure rate.

Limit / assumption

The formula omits logistics delay, preventive maintenance, common-cause failures, waiting for crew access and multiple repair states unless those are folded into MTTR.

Formula 2 — spare demand from expected failures

Starting question

How many failures should planners roughly expect from a fleet over a declared operating exposure?

Read aloud

Read: “expected failures equal total fleet operating hours divided by mean time between failures.”

Symbols, pronunciation and meaning

H_{fleet} is the sum of operating hours across the relevant units; MTBF is the average hours between failures for the component or item.

Units

Both quantities use hours, so the quotient is an expected count of failures, not a guaranteed integer outcome.

Origin and status of values

Fleet hours come from the mission plan or logs; MTBF must correspond to the same component, duty cycle and environment as closely as possible.

Why this operation

If one failure occurs on average per MTBF hours, dividing total exposure by that interval gives the statistical expectation.

Substitution and calculation

Twelve units each operating 2,000 h create 24,000 fleet-hours. With MTBF 6,000 h, expected failures $\approx 24,000 / 6,000 = 4$.

Calculator entry

Calculate fleet hours first: 12×2000 . Then divide by 6000. Keep the assumptions visible.

Mental estimate

Each 6,000 hours contributes about one expected failure; 24,000 hours therefore suggests roughly four.

Independent check

Multiply four expected failures by 6,000 h and recover the 24,000 h exposure.

Physical or operational interpretation

Four expected failures does not mean four spares are automatically sufficient; variability, repairability, common failures and resupply delay matter.

Plain-English translation

The planning model says this fleet exposure is of the order of four failures for the stated reliability.

Variation / sensitivity

If duty doubles, expected failures double. If MTBF improves by 50%, expected demand falls proportionally in this simple model.

Limit / assumption

Failure processes may not be constant-rate, especially for wear-out, early-life defects and condition-dependent damage. Use distributions when evidence supports them.

Mission reasoning: restore function, then improve the fleet

Design access before launch

A component that can theoretically be replaced may still be unmaintainable if a crew must remove ten unrelated assemblies, vent a habitat zone or use a tool that cannot fit. Maintainability reviews should inspect access path, connectors, lifting, contamination control, fasteners and the ability to perform the task in gloves or other mission constraints.

Diagnose before swapping parts

Blind replacement consumes spares and can leave the real fault untouched. Maintenance procedures should help crews separate symptom from cause using measurements, built-in test, substitution and isolation. A fault tree or troubleshooting logic is especially valuable when several subsystems can produce the same observed symptom.

Treat tools as spares enablers

A stockroom full of components does not create repair capability if the settlement lacks pullers, torque tools, seals, lubricants, test equipment or calibration. Tool availability should be linked to the maintenance task list. Rare tools may justify redundancy when a single loss could strand many repairs.

Control cannibalization explicitly

Removing a working assembly from a parked rover can restore a critical rover quickly, but the donor vehicle must be reclassified, its missing configuration recorded and the borrowed part traced. Informal cannibalization creates “ghost defects” in the fleet and makes inventory records untrustworthy.

Close the loop from failure to redesign

Repeated maintenance should change engineering. If one connector fails every dusty season or one filter clogs twice as fast as expected, the settlement should not merely stock more of it forever. Data can justify shielding, procedure changes, new materials or local redesign that permanently reduces maintenance burden.

Maintenance exercises — reason from failure data to repair strategy

Exercise A — Availability trade

A pump has MTBF 800 h and MTTR 40 h. Compute simple availability. Then estimate the effect of reducing MTTR to 10 h.

Reveal the reasoned solution

Initial availability = $800 / (800 + 40) \approx 0.9524$, or 95.2%. With a 10 h MTTR, availability = $800 / (800 + 10) \approx 98.8\%$. The example shows why better access, tools and diagnostics can improve service availability even if the pump itself is no more reliable.

Exercise B — Fleet exposure

Eight identical fans each run 3,000 h. Their planning MTBF is 4,000 h. Estimate expected failures.

Reveal the reasoned solution

Fleet exposure is $8 \times 3,000 = 24,000$ h. Expected failures are about $24,000 / 4,000 = 6$. This is a statistical planning figure; spare quantity should account for uncertainty, repair of failed fans and common-cause risk.

Exercise C — Single-tool dependency

A special puller is required to replace bearings on six critical machines, and only one puller exists. Classify the risk.

Reveal the reasoned solution

The puller is a common maintenance dependency. Losing it can simultaneously remove the repairability of six otherwise independent machines. Options include a redundant puller, a locally manufacturable design, alternate removal method or redesign of the bearing interface.

Exercise D — Cannibalization decision

One of four rovers is already unavailable for a motor fault. A second rover needs its compatible navigation computer for a rescue task. What must be recorded if the computer is borrowed?

Reveal the reasoned solution

Record donor and receiver identifiers, part serial or lot, removal and installation condition, software/configuration state, tests performed, the donor rover’s new status, and the plan to restore the donor. Without this, the fleet inventory becomes fiction.

Exercise E — Repeated seal failure

Three identical seals fail after dusty EVAs. What evidence would justify changing the design rather than only stocking more seals?

Reveal the reasoned solution

Look for a repeatable relation between dust exposure, seal location, inspection findings and failure mode; verify installation was correct; compare lots; inspect surface wear; and test proposed shielding or material changes. A redesign should target demonstrated cause, not simply correlation.

Exercise F — Repair queue

Two failed units each need six hours of hands-on work, but only one qualified technician is available for four hours per day. What is the minimum calendar time for the hands-on portion?

Reveal the reasoned solution

Total hands-on work is 12 hours. At four technician-hours per day, at least three days are required before adding diagnosis, waiting, testing or interruptions. Crew skill availability can therefore be a maintenance bottleneck just like spare parts.

Interactive beginner glossary

The following terms let a crew discuss reliability, repair and fleet status precisely instead of using the word “broken” for every maintenance condition.

- failure — Loss of the ability of an item to perform a required function within its stated limits.
- fault — An abnormal condition or defect that can contribute to a failure; a fault may exist before the required function is lost.
- failure mode — The observable way in which a component or function fails, such as leaking, seizing, drifting or losing communication.
- failure cause — The physical, procedural, software or environmental mechanism that produced a failure mode.
- MTBF — Mean Time Between Failures, an average operating interval between failures for repairable equipment under stated conditions.
- MTTR — Mean Time To Repair, an average time required to restore failed equipment under stated maintenance conditions.
- availability — Fraction of the required time during which a system or item is capable of performing its required function.
- maintainability — A design property describing how effectively an item can be inspected, serviced, diagnosed, removed, repaired and restored.
- preventive maintenance — Scheduled maintenance performed before known failure, often based on time, cycles or usage.
- predictive maintenance — Maintenance initiated from measured condition or trends intended to indicate developing degradation before functional failure.
- corrective maintenance — Maintenance performed to diagnose and restore an item after a fault or failure has been detected.
- condition monitoring — Repeated measurement of indicators such as vibration, temperature, leakage or electrical behaviour to detect degradation.
- built-in test — Diagnostic capability embedded in equipment to detect, isolate or report faults.
- troubleshooting — Structured process of using symptoms and tests to isolate the cause of a problem.
- replaceable unit — A component or assembly deliberately defined as a unit that can be removed and replaced during maintenance.
- spare — A replacement item held in inventory to restore capability after failure, wear or planned maintenance.
- rotatable — A repairable spare that cycles between installed service, removal, repair and return to inventory.
- consumable — An item that is used up or discarded during operations or maintenance and is not normally restored for reuse.
- cannibalization — Removal of a serviceable part from one asset to restore another asset.
- configuration control — Management of the documented hardware, software and procedural state so operators know what each asset actually contains.
- maintenance log — Record of symptoms, diagnosis, actions, parts, measurements and return-to-service evidence for maintenance events.

- work order — Controlled authorization and record for a maintenance task, including scope, resources, status and completion evidence.
- torque — Twisting moment applied to a fastener or shaft; correct torque can be essential to joint integrity.
- lubricant — Material used to reduce friction or wear between moving surfaces and sometimes to provide sealing or corrosion protection.
- seal — Component or interface intended to prevent unwanted leakage of gas, liquid, dust or contaminants.
- wear-out — Failure tendency that increases as accumulated cycles, friction, fatigue or material degradation approaches a life limit.
- common-cause failure — A single cause or condition that can defeat multiple items that were assumed to provide independent redundancy.
- mean logistics delay — Average time lost waiting for parts, tools, access or other support before hands-on repair can proceed.
- return to service — Formal decision that maintenance is complete and the repaired item has passed required checks for operational use.
- reliability growth — Improvement of reliability through testing, failure analysis, corrective action and redesign rather than by assumption alone.

Operational depth: turn the workshop into a reliability-learning system

Separate urgent restoration from root-cause work

A crew may need to restore a life-support pump immediately using a known spare. That does not end the engineering task. The removed unit should later be examined so the settlement can determine whether the fault was wear, contamination, installation, software or a common design weakness.

Protect repair time from poor access

Maintenance burden often comes from the hours around the actual replacement: depressurizing, moving panels, cleaning dust, fetching tools and retesting. Recording those steps separately reveals whether redesigning access would save more crew time than buying a slightly more reliable component.

Use fleet data carefully

A tiny Mars fleet produces sparse statistics. Engineers should combine observed failures with physical inspection and appropriate Earth test data rather than claiming precise reliability from a handful of events. Confidence bounds and engineering judgement matter when data are limited.

Plan for repaired spares, not only new spares

If a failed module can be diagnosed and restored locally, one spare can support several failure cycles. The spares model should therefore include repair turn-around time, test capacity and parts consumed inside the repair. A nominally repairable item that cannot be tested after repair is operationally questionable.

Control maintenance documentation like software

Procedures evolve after lessons and design changes. The crew must know which revision applies to each configuration. An old procedure used on new hardware can create a new failure while appearing compliant. Revision status should therefore be visible at the point of work.

Reward reporting of near failures

Intermittent noise, rising motor current or a connector that nearly loosens can be more valuable than a clean success report. A culture that hides small anomalies to protect schedule loses the chance to intervene before the failure becomes mission-threatening.

Measure the maintenance burden by function, not only by part count

A small number of high-burden tasks can dominate crew maintenance time. One filter replacement that requires suit work, depressurization and an hour of cleaning may cost more operationally than dozens of simple electronic swaps. The fleet database should therefore record hands-on time, access time, waiting time and test time separately. That information helps designers eliminate the expensive steps and prevents spare-count statistics from hiding the true cost of keeping a function available.

Operational review checklist

- List required maintenance tasks while the hardware is still being designed.
- Check physical access, tools, contamination control and post-repair testing for each critical replaceable unit.
- Distinguish symptom, failure mode and root cause in maintenance records.
- Estimate availability using both failure frequency and restoration time.
- Size spares from fleet exposure, uncertainty, repairability and resupply delay.
- Treat specialist tools and test equipment as part of the spare strategy.
- Record every cannibalized part and donor configuration.
- Keep maintenance procedures under revision and configuration control.
- Analyse repeated failures for design or procedure changes, not only inventory increases.
- Use repaired-item evidence to update reliability assumptions and future mission planning.

Primary sources and bridges

- NASA — Reliability and Maintainability
- NASA-STD-8729.1 — Reliability and Maintainability
- NASA NTRS — Supportability Concepts for Crewed Deep Space Exploration

FR ↔ EN parity matrix

Intrinsic availability	Disponibilité intrinsèque
Expected failures	Défaillances attendues
Maintenance workload	Charge de maintenance
Spare-stock coverage	Couverture du stock de rechanges
Fleet availability	Disponibilité de flotte
Series-chain reliability	Fiabilité d'une chaîne série