

Diagnostics and FDIR: detecting, isolating and recovering from anomalies

Guiding question

Concrete → interfaces → measurement → calculation → failure → recovery.

Course map

The real phenomenon → Vocabulary and problem boundary → Course-specific system view → Mathematical relationship and reading the symbols → Worked calculations and interpretation

1 — The real phenomenon

FDIR means Fault Detection, Isolation and Recovery. Detecting an anomaly is not enough: a bad sensor, degraded equipment, common cause and secondary effect must be distinguished before selecting a safe recovery. Observability depends on sensor quality and consistency models.

The guiding question is: How do we know what actually failed when several faults produce similar symptoms? Reasoning starts with the physical or operational function before introducing the mathematical relationship. The goal is not to accumulate terminology, but to know which quantity changes, why it changes and what becomes hazardous when it leaves its domain. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the first task here is therefore to identify the mechanism specific to this subject before searching for an equation or reference value.

2 — Vocabulary and problem boundary

In “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, distinguish the phenomenon, available measurement, any command, the margin and the success criterion. The calculation boundary states what is included and excluded; without that boundary, a percentage, mass or time may be mathematically correct but wrong as an engineering conclusion. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the chosen boundary also states what would otherwise be double-counted or omitted from a mission budget.

memory errors, resets, dose, current, temperature, sensor availability, data consistency and event logs bit upset, misconfigured watchdog, bad time reference or redundant channels sharing the same supply error injection, irradiation where relevant, software tests, data verification and safe-mode simulation

Equation - evidence - decision

Résidu = mesure – prédiction

Failure to test: The reference failure is not a vague “broken component.” For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, test in particular bit upset, misconfigured watchdog, bad time reference or redundant channels sharing the same supply. Diagnosis asks which symptoms appear first, which are only consequences and which action preserves the most options.

Diagnostics and FDIR: detecting, isolating and recovering from anomalies

3 — Course-specific system view

This lesson does not reuse one generic picture for every subject. The system view follows cause → measured quantity → decision or physical response → limit for “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”. The English text remains fully equivalent while large translated illustrations are intentionally deferred until their dedicated artwork is supplied. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the system view must expose inputs, outputs, measured quantity and the consequence of drift without relying on a generic module diagram.

4 — Mathematical relationship and reading the symbols

Read aloud : a residual is the difference between a measurement and the model prediction.

Before substituting numbers, write the unit of every term, state whether the relationship is a physical law, approximation or project indicator, and check dimensional consistency. This is especially important here because “Diagnostics and FDIR: detecting, isolating and recovering from anomalies” combines quantities that do not all have the same evidence status. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, this relationship is chosen because of the phenomenon under study; a different dominant quantity would require a different equation or model.

5 — Worked calculations and interpretation

$102.4 - 100.0 = +2.4$ units

$|2.4| > 2.0 \rightarrow$ alarm in this example

Two sensors at 50 and 51; one at 88 → third is suspect, to be confirmed

- 1. 1. Residual $102.4 - 100.0 = +2.4$ units Interpretation: this result is used only after comparison with units, margin and the scenario boundary for “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”.
- 2. 2. Threshold $|2.4| > 2.0 \rightarrow$ alarm in this example Interpretation: this result is used only after comparison with units, margin and the scenario boundary for “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”.
- 3. 3. Sensor vote Two sensors at 50 and 51; one at 88 → third is suspect, to be confirmed Interpretation: this result is used only after comparison with units, margin and the scenario boundary for “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”.

Operational table

Phenomena / sections:	3 — Course-specific system view ; 4 — Mathematical relationship and reading the symbols ; 5 — Worked calculations and interpretation
Relationship to revisit:	Résidu = mesure – prédiction
Expected decision:	measure → compare with margin → choose nominal, degraded or stop

Check before continuing

Explain without rereading: 3 — Course-specific system view ; 4 — Mathematical relationship and reading the symbols ; 5 — Worked calculations and interpretation. State one measured quantity, one margin and one possible decision.

Diagnostics and FDIR: detecting, isolating and recovering from anomalies

6 — What the formula does not contain

The relationship “Résidu = mesure – prédiction” does not by itself contain all of “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”. It does not automatically tell us whether a sensor is valid, a structure is aging, a resource is accessible, a command arrives in time or a secondary failure removes margin. The example $102.4 - 100.0 = +2.4$ units therefore remains a local calculation rather than a complete architecture.

To make the model useful, explicitly add the quantities that dominate this subject: memory errors, resets, dose, current, temperature, sensor availability, data consistency and event logs. We can then ask which variation truly changes the result, which is negligible and which forces an architectural change. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, this model limitation states exactly what a correct calculation still cannot establish about the real system.

7 — Instrumentation, observability and data quality

For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, observability relies on memory errors, resets, dose, current, temperature, sensor availability, data consistency and event logs. Each datum has a unit, acquisition rate, uncertainty, timestamp and validity domain. A value arriving without context can be more dangerous than no measurement because it creates unjustified confidence. Consistency is checked with at least one independent piece of information when the function is critical. A trend, physical balance or second measurement principle helps distinguish a real system change from a drifting sensor. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the selected instrumentation must distinguish a real physical change from sensor drift or a bad state estimate.

8 — Phenomenon-specific failures and recovery

The reference failure is not a vague “broken component.” For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, test in particular bit upset, misconfigured watchdog, bad time reference or redundant channels sharing the same supply. Diagnosis asks which symptoms appear first, which are only consequences and which action preserves the most options.

The degraded mode must be defined before failure: minimum function, allowable duration, consumed stock, crew action, abort condition and return-to-nominal criterion. That sequence is topic-specific and cannot be replaced by one universal paragraph about redundancy. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the degraded mode is defined around the minimum function specific to this subject, with an abort threshold and a return-to-nominal condition.

Operational table

Phenomena / sections:	6 — What the formula does not contain ; 7 — Instrumentation, observability and data quality ; 8 — Phenomenon-specific failures and recovery
Relationship to revisit:	Résidu = mesure – prédiction
Expected decision:	measure → compare with margin → choose nominal, degraded or stop

Check before continuing

Explain without rereading: 6 — What the formula does not contain ; 7 — Instrumentation, observability and data quality ; 8 — Phenomenon-specific failures and recovery. State one measured quantity, one margin and one possible decision.

Diagnostics and FDIR: detecting, isolating and recovering from anomalies

9 — NASA / reference case

NASA material is used as an evidence dossier: requirements, reliability, maintainability, testing and configuration. The lesson never turns a generic failure rate into a universal truth; it shows how evidence is bounded to defined hardware, environment and duration.

The case is used only within what it actually demonstrates. Flight measurement, human-system standard, component test and architecture study are different kinds of evidence; the text therefore states what is observed, calculated, simulated or still prospective. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the cited NASA case is used as targeted evidence for this phenomenon and is never turned into one universal Mars architecture.

10 — Architecture trade

A good solution for “Diagnostics and FDIR: detecting, isolating and recovering from anomalies” does not maximize one metric. Compare nominal performance, mass, energy, simplicity, maintenance, crew time, common dependencies and recoverability. An option that improves $[2.4] > 2.0 \rightarrow$ alarm in this example can still be rejected if it makes failure detection or repair much harder.

The trade is recorded together with its assumptions. If environment data, mass or mission cadence changes, we know which conclusions must be recomputed instead of silently preserving an obsolete choice. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the trade is evaluated against the interfaces actually touched by this subject rather than a generic list of desirable qualities.

11 — Demonstration, testing and success criteria

The evidence strategy for “Diagnostics and FDIR: detecting, isolating and recovering from anomalies” combines error injection, irradiation where relevant, software tests, data verification and safe-mode simulation. Every test records exact hardware, software, configuration, environment, tolerances and success criterion. A successful demonstration outside the mission domain does not replace qualification inside it.

Evidence grows by levels: analytical relationship, simulation, component, subsystem, integrated system, duration and failure. This hierarchy prevents one spectacular test from being presented as validation of the whole mission. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, demonstration must reproduce the constraints that make this phenomenon difficult; a spectacular test outside the mission domain is insufficient.

Operational table

Phenomena / sections:	9 — NASA / reference case ; 10 — Architecture trade ; 11 — Demonstration, testing and success criteria
Relationship to revisit:	Résidu = mesure – prédiction
Expected decision:	measure $\rightarrow$ compare with margin $\rightarrow$ choose nominal, degraded or stop

Check before continuing

Explain without rereading: 9 — NASA / reference case ; 10 — Architecture trade ; 11 — Demonstration, testing and success criteria. State one measured quantity, one margin and one possible decision.

Diagnostics and FDIR: detecting, isolating and recovering from anomalies

12 — Decision exercise

Situation: revisit “Diagnostics and FDIR: detecting, isolating and recovering from anomalies” with a 20% increase in the most penalizing quantity from the first worked example while one measurement or backup path is unavailable.

Situation: revisit “Diagnostics and FDIR: detecting, isolating and recovering from anomalies” with a 20% increase in the most penalizing quantity from the first worked example while one measurement or backup path is unavailable. Expected answer: recompute the relationship, identify remaining margin, check whether observability is still adequate, and decide whether degraded operation remains acceptable. Multiplying by 1.2 is not enough if the variation also changes interfaces or limits.

13 — What to retain without over-generalizing

14 — Topic-specific primary sources

These references directly document the phenomenon, technology or human constraint addressed in this lesson. They do not by themselves define an official Mars architecture. For “Diagnostics and FDIR: detecting, isolating and recovering from anomalies”, the bibliography is deliberately targeted to this page so that readers can trace each claim back to the relevant primary document.

Topic-specific primary sources

NASA NTRS — Lessons Learned on Implementing Fault Detection, Isolation and Recovery - <https://ntrs.nasa.gov/citations/20130012528>

NASA NTRS — FDIR for Autonomous Space Systems - <https://ntrs.nasa.gov/citations/20250003745>

NASA-STD-8729.1A — Reliability and Maintainability Standard - <https://standards.nasa.gov/standard/NASA/NASA-STD-87291>

Operational table

Phenomena / sections:	12 — Decision exercise ; 13 — What to retain without over-generalizing ; 14 — Topic-specific primary sources
Relationship to revisit:	Résidu = mesure – prédiction
Expected decision:	measure → compare with margin → choose nominal, degraded or stop

Check before continuing

Explain without rereading: 12 — Decision exercise ; 13 — What to retain without over-generalizing ; 14 — Topic-specific primary sources. State one measured quantity, one margin and one possible decision.