PE

PROFESSIONAL ENGINEERING CONTINUING EDUCATION

Engineering Reliability & Maintainability

Failure Prevention, Redundancy & Asset Performance

Self-Paced Course Study Guide

DEFINE SUCCESS

Function, duty, mission and operating context.

→

DESIGN / SUPPORT

Independent paths, safe access and feasible restoration.

→

LEARN / VERIFY

Operating evidence, corrective action and recurrence checks.

Reliability decisions connect required function, design and support, operating evidence and verified improvement.

COURSE	CREDIT / CATEGORY	INSTRUCTOR
ENG-012	2 CEH / PDH · Pertinent to Engineering Professional Practice / Reliability & Asset Management	Matthew Ward, MBA, CIA

Superior Learning, LLC d/b/a BuilderCE

For professional engineers and professionals responsible for engineered assets.

ENGINEERING RELIABILITY & MAINTAINABILITY

2 CEH

PERTINENT

ENGINEERING

Course Overview

Course purpose

Reliability is the ability of an engineered item or system to perform its required function for the required period and under the stated conditions. Maintainability is the ability to retain or restore that function using defined personnel, procedures, tools, spares, access, and support. Availability connects the two: an asset cannot deliver value merely because it rarely fails if each failure takes days to diagnose and repair; nor is fast repair sufficient if failures are frequent, unsafe, or disruptive. The engineering objective is dependable performance across the asset life cycle, not simply a lower failure count.

This course develops a practical framework for professional engineers who influence design, procurement, operation, maintenance, renewal, or technical oversight. It emphasizes the decisions that make reliability and maintainability tangible: defining required functions and operating context; distinguishing failure mode from failure effect; using MTBF, MTTF, MTTR, availability, and failure-rate concepts correctly; understanding series and parallel arrangements; recognizing common-cause dependencies; designing for access and diagnostics; selecting preventive, condition-based, failure-finding, run-to-failure, or redesign actions; and using operating data to improve the asset rather than merely document its history.

The course is intentionally cross-disciplinary. A pump train, emergency generator, bridge movable system, building automation network, wastewater blower, transportation signal cabinet, industrial process skid, and mission-support test system have different technical details, but the reliability logic is transferable. Each has required functions, failure modes, consequences, detection opportunities, restoration tasks, dependencies, and evidence that can be used to improve performance.

Reliability-centered maintenance (RCM) is addressed as a structured method for selecting failure-management strategies, not as a slogan for “more preventive maintenance.” Current public guidance from NASA and the U.S. Department of Energy emphasizes matching maintenance to operating context, function, failure behavior, consequence, and cost-effectiveness. SAE JA1011 provides evaluation criteria for an RCM process, while ISO 14224 provides a standardized reliability-and-maintenance data framework for petroleum, petrochemical, and natural-gas equipment. Those sources inform the course framework; project-specific codes, owner requirements, manufacturer instructions, and discipline standards remain controlling where applicable.

Learning objectives

- Explain the distinct roles of reliability, maintainability, availability, supportability, and asset criticality in engineering decisions.
- Apply basic reliability concepts including failure rate, MTBF/MTTF, MTTR, reliability over time, series/parallel arrangements, redundancy, and common-cause failure.
- Evaluate design and procurement choices for maintainability, including access, isolation, diagnostics, modularity, standardization, documentation, and spare-part support.
- Select among reactive, time-based preventive, condition-based, failure-finding, and redesign strategies based on failure behavior and consequence rather than habit.
- Use reliability-centered maintenance logic to connect required function, functional failure, failure mode, consequence, and technically effective maintenance actions.
- Interpret operating and maintenance data, distinguish useful metrics from misleading averages, and use recurring-failure evidence to drive reliability growth and asset improvement.

CONCEPT	ENGINEERING QUESTION
Reliability	Will the required function continue without failure over the stated mission and conditions?
Availability	Is the required function ready when needed, accounting for the defined failure and restoration clocks?
Maintainability	Can function be restored safely within the required time using the prescribed resources?

Study plan and navigation

MODULE	MINUTES	FOCUS
1	25	Foundations of Reliability, Maintainability & Availability
2	25	Reliability Modeling, Redundancy & Failure Mechanisms
3	25	Maintainability, Maintenance Strategy & Condition Monitoring
4	25	Reliability-Centered Maintenance, Criticality & Asset Strategy
5	20	Reliability Data, Growth & Applied Engineering Decisions

Planned instructional time: 120 minutes, exclusive of the final examination. Individual study time may vary.

How to complete this course

Read all modules in this course. No separate module acknowledgement is required. When you have finished reading, select Continue at the bottom of the BuilderCE lesson page to begin the 10-question final examination. If you are reading the PDF, return to that webpage first. A score of at least 70% is required to pass. Retakes are permitted.

Thermal evidence and completed diagnostic record · Sources

> **Course boundary.** This course teaches general reliability, maintainability, and asset-performance concepts. It does not replace manufacturer maintenance instructions, mandatory inspection intervals, code requirements, reliability demonstration requirements, safety analyses, discipline-specific design standards, or the judgment of an appropriately qualified engineer.

MODULE 1 | 25 MINUTES

Foundations of Reliability, Maintainability & Availability

1.1 Required function comes before the failure metric

A reliability number has meaning only in relation to a required function and operating context. “The pump is reliable” is incomplete unless the engineer knows the required flow, head, fluid, duty cycle, operating environment, mission time, allowable degradation, and what constitutes failure. A standby fire pump that starts successfully on test but cannot deliver the required pressure is not functionally reliable. A valve that passes flow but cannot reach its required fail-safe position has lost a different function. A bridge movable-span drive may have separate functions for movement, holding, braking, position indication, and emergency recovery.

This functional viewpoint prevents a common mistake: counting component failures without asking whether the system still delivered the required performance. Reliability analysis should begin by defining function in measurable terms wherever practical. Capacity, accuracy, response time, leakage, vibration, power quality, environmental range, structural limit, service availability, or another performance measure may define success. The definition should also identify the operating context because the same asset can have different failure behavior under continuous duty, intermittent standby, corrosive exposure, vibration, cycling, or temperature extremes.

Functional failure can be total or partial. A chiller that can produce only 70 percent of rated capacity during peak conditions has not “failed” in the same sense as a seized compressor, but it may still fail the facility requirement. Partial degradation is especially important for condition monitoring because the asset may be operating while its margin is shrinking. Engineers should distinguish observable symptoms, degraded states, and loss of required function so that detection thresholds are connected to useful action.

1.2 Reliability, availability, and maintainability are related but different

Reliability describes the probability that an item performs without failure over a stated period under stated conditions. Maintainability describes the ease and probability of retaining or restoring the item within a specified time when maintenance is performed with prescribed procedures and resources. Availability describes readiness for use and is influenced by both how often failures occur and how long restoration takes. These distinctions are reflected in NASA reliability and maintainability guidance and in common engineering practice.

Consider two identical process trains. Train A fails twice per year and requires forty hours for each repair. Train B fails four times per year but is restored in two hours each time. Train A has better failure frequency but may have worse availability. Which is preferable depends on consequence, redundancy, production requirements, safety, maintenance staffing, and the ability to schedule work. This example shows why “improve MTBF” is not always the right objective. Sometimes reducing MTTR or eliminating a long-lead spare is the faster path to improved service.

Maintainability is not merely a maintenance-department responsibility. Access platforms, lifting points, isolation valves, connector placement, cable routing, equipment spacing, diagnostic coverage, software architecture, modularity, fastener selection, labeling, configuration control, and documentation are design decisions. By the time an asset is operating, many maintainability opportunities have already been fixed by geometry and architecture.

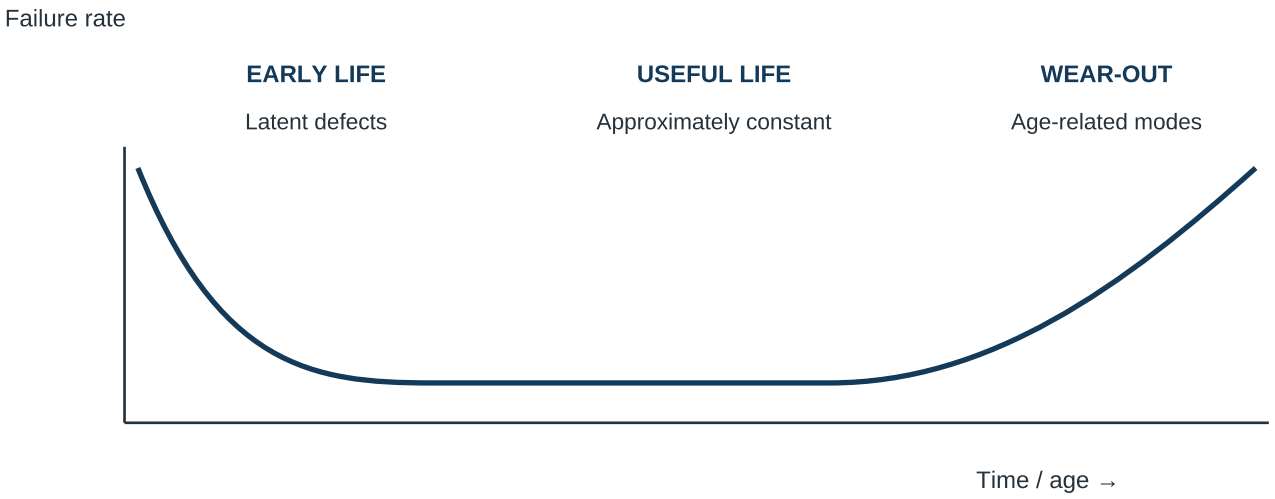
SIMPLE STEADY-STATE MODEL	INTERPRETATION
A = MTBF / (MTBF + MTTR)	Applicable to a stable alternating operating / restoration process with consistent clock definitions.
Train A: 2 failures × 40 h = 80 h/year	Lower failure count, greater annual outage time in this example.
Train B: 4 failures × 2 h = 8 h/year	More failures, shorter total outage time. Consequence still determines the preferred option.

1.3 Failure rate, MTBF, and MTTF: useful but easy to misuse

Failure rate is a measure of how frequently failures occur relative to exposure. Exposure might be operating hours, cycles, starts, miles, demands, or another basis appropriate to the asset. Mean time between failures (MTBF) is commonly used for repairable items; mean time to failure (MTTF) is commonly used for nonrepairable items. These are averages over a defined population and condition. They are not guaranteed life values.

An MTBF of 10,000 hours does not mean an individual component will run for 10,000 hours and then fail. Under a constant-hazard exponential model, a 10,000-hour MTBF corresponds to a reliability of about 36.8 percent at 10,000 hours. This surprises people because they interpret the mean as a “life expectancy” threshold. The model also assumes a constant hazard rate, which may be inappropriate for wear-out, infant mortality, or failure mechanisms strongly dependent on age, usage, environment, or load.

The exposure basis must match the mechanism. A relay that operates only on demand may be better described by failures per demand. A bearing may be related to operating hours and load. A fatigue-sensitive component may be related to cycles. A battery may be influenced by calendar age, temperature, cycling, and depth of discharge. Engineers should resist forcing every asset into an hours-based metric simply because a maintenance system reports hours conveniently.



Qualitative lifecycle failure-rate model. The shape is not measured fleet data; actual mechanisms and mixed populations may behave differently.

1.4 The bathtub curve is a model, not a law

The classic bathtub curve combines three conceptual regions: decreasing early-life failure rate, an approximately constant useful-life region, and increasing wear-out failure rate. It is useful for explaining why different failure populations require different strategies. Early-life failures may be associated with manufacturing defects, installation errors, commissioning weaknesses, contamination, or latent defects. Useful-life failures may appear more random. Wear-out may reflect fatigue, corrosion, insulation aging, erosion, seal degradation, wear, embrittlement, or another accumulating mechanism.

However, real assets are mixtures of components and failure modes. A pump can simultaneously experience random electrical failures, age-related seal wear, condition-dependent cavitation damage, and maintenance-induced infant mortality after overhaul. A complex system therefore rarely has a single clean bathtub curve. The engineering value of the concept is to ask whether failure likelihood is meaningfully related to age or usage for the specific failure mode being managed.

This distinction affects preventive maintenance. Replacing a component on a fixed calendar interval is technically justified when the failure mode has a sufficiently predictable age or usage relationship and the task restores resistance to that failure. If failures are dominated by random mechanisms, periodic replacement can add cost and may even introduce installation-related defects without materially reducing risk.

1.5 Failure, defect, degradation, and cause

A defect is a nonconforming condition; degradation is a change in condition or performance; failure is loss of required function. A defect may exist without immediate functional failure. Degradation may be measurable long before performance falls below the required threshold. The failure cause is the

physical, human, environmental, design, manufacturing, installation, maintenance, or organizational factor that produced the failure mode. Using these terms consistently improves analysis and data quality.

For example, elevated bearing vibration is a symptom, not necessarily the failure mode. Misalignment may be a cause, excessive vibration a degraded condition, bearing damage a failure mechanism, and inability to maintain shaft speed a functional failure. Treating every alarm as a “failure” and every failure as a generic “bearing problem” destroys the information needed for recurrence prevention.

Good reliability records therefore separate what was observed from what was concluded. Initial work orders can record symptoms and provisional codes. After technical review, the record can be updated with confirmed failure mode and cause where evidence supports it. This is more useful than forcing technicians to guess root cause at the moment of response.

MODULE 2 | 25 MINUTES

Reliability Modeling, Redundancy & Failure Mechanisms

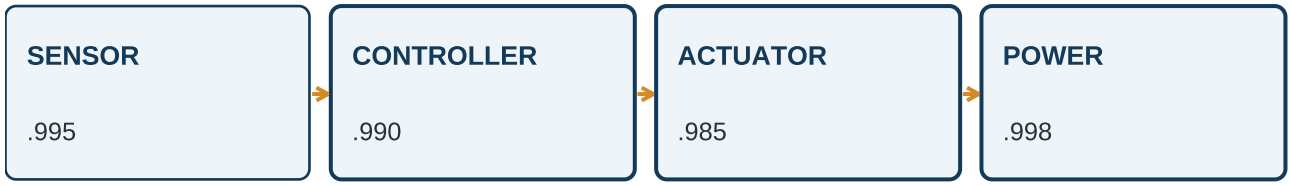
2.1 Series systems: every required element matters

In a pure series reliability model, the system succeeds only if every required component succeeds. If component failures are independent and each component reliability over the mission is known, system reliability is the product of the component reliabilities. This mathematical structure captures an important engineering reality: adding required components can reduce system reliability even when every component is individually “high reliability.”

Suppose four required elements each have reliability of 0.98 over a mission interval. The series-system reliability is approximately 0.922. Ten such required elements produce about 0.817. This does not mean complex systems are inherently unacceptable; it means architecture, failure tolerance, diagnostics, maintenance, and redundancy become more important as the number of required elements grows.

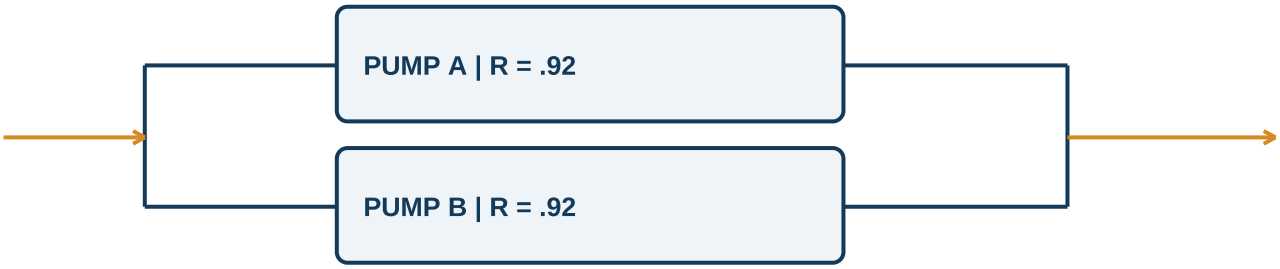
Series models also help identify weak links, but the engineer must verify that the logical model matches the physical function. Some components are bypassable, some failures are degraded rather than total, and some functions have multiple success paths. Reliability block diagrams are functional models, not necessarily wiring diagrams or process-flow diagrams.

SERIES: every required element must succeed



R = .995 × .990 × .985 × .998 = .9683

PARALLEL: either independent full-capacity path succeeds



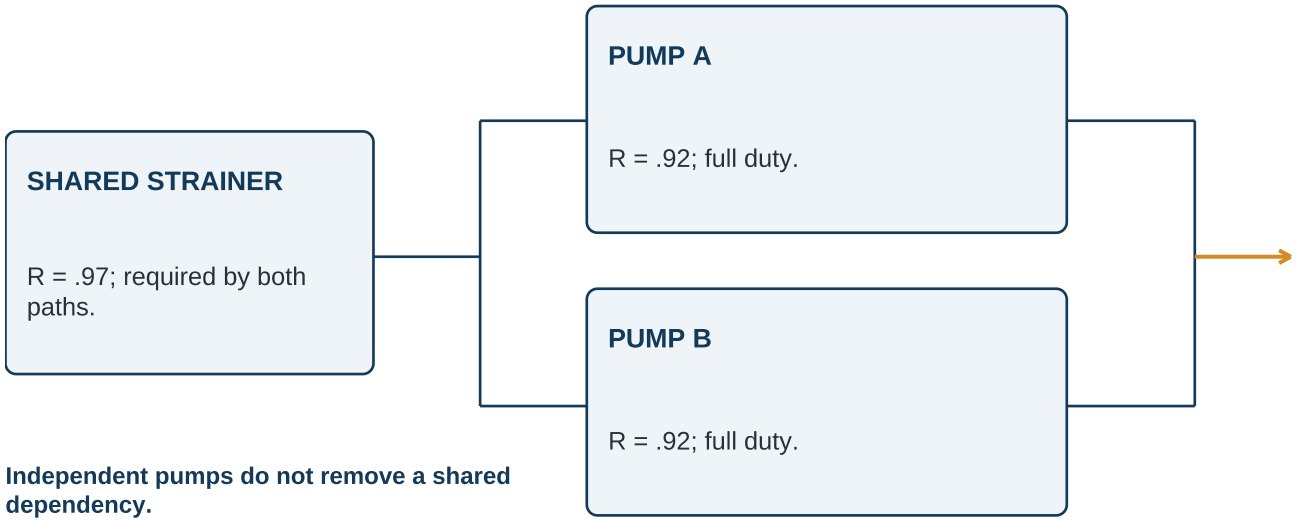
Independent mission models. Series success requires every element; ideal parallel success requires either full-capacity path. Repair and switching behavior are not modeled.

2.2 Parallel redundancy can improve reliability - when the paths are truly independent

Parallel paths can improve reliability because the system can succeed when one path fails. For two identical independent components with reliability R arranged so either can perform the required function, system reliability is 1 - (1-R)^2. If each has reliability 0.90 over the mission, the idealized parallel system has reliability 0.99. The apparent improvement is substantial.

That improvement depends on assumptions that are often violated. The two components must each have adequate capacity; the system must detect and isolate failures correctly; switching or load-sharing logic must work; shared support systems must remain available; and the failures must be sufficiently independent. If both pumps share suction blockage, both generators share fuel contamination, both controllers share corrupted software, or both communication paths use the same conduit, the nominal redundancy can be defeated by one event.

Redundancy can also increase complexity. Additional valves, controls, breakers, software, sensors, and maintenance tasks can create new failure opportunities. Standby equipment may degrade silently because it is not exercised. Active load-sharing equipment can accumulate different wear than a cold standby. Engineers should evaluate the reliability of the complete success path, not merely count redundant boxes.



Shared strainer reliability multiplies the ideal parallel reliability only when the support and pump failures are independent. Common environmental or operational causes require additional analysis.

Common-cause failure occurs when multiple components fail because of a shared cause or dependency. Common-mode failure is often used similarly, although terminology can vary by industry. Examples include flooding multiple electrical rooms, a single software defect affecting redundant controllers, shared cooling loss, common maintenance error, contaminated lubricant supplied to multiple machines, a design error repeated across identical units, or a fire damaging both redundant cable routes.

The engineering response begins with dependency mapping. Ask what the redundant paths share: power, cooling, fuel, network, physical space, environment, human action, maintenance procedure, supplier, software, calibration standard, or upstream process. The system may have independent equipment but still share a hidden support dependency. Independence should be demonstrated at the level of the failure mechanisms that matter to the required function.

Separation and diversity can reduce common-cause exposure. Physical separation protects against localized hazards; electrical separation protects against common supply faults; diverse technologies may reduce common design or software failure; separate maintenance teams or independent verification may reduce repeated human error. Diversity is not automatically better, because it can increase training and spare complexity. The appropriate balance depends on consequence and evidence.

2.4 Reliability functions and the constant-hazard model

When a constant failure rate is a reasonable approximation, reliability over time can be modeled as R(t) = exp(-lambda t), where lambda is the failure rate. MTBF is the reciprocal of lambda for the exponential model. This relationship is useful for mission reliability, spare planning, and comparing architectures, but only when the assumption is appropriate.

If lambda is 1 failure per 20,000 operating hours, the probability of surviving 1,000 hours without failure under the model is exp(-1000/20000), or about 95.1 percent. Survival over 10,000 hours is about 60.7 percent. The same underlying rate produces different reliability depending on the mission duration, which is why a reliability percentage without a time basis is incomplete.

The exponential model is attractive because it is simple and memoryless: conditional on surviving to the present, the future hazard is unchanged. Real degradation mechanisms often violate that assumption. Engineers should use simple models as screening tools and escalate to age-dependent or condition-dependent models when the decision is sensitive to the assumption.

2.5 Weibull thinking and age-related failure behavior

The Weibull distribution is widely used because its shape parameter can represent decreasing, roughly constant, or increasing hazard behavior. A shape parameter beta less than one is associated with decreasing hazard; beta near one resembles constant hazard; beta greater than one represents increasing hazard with age or usage. The scale parameter sets the time scale.

Weibull analysis can help determine whether a time-based replacement policy has a technical basis, but it should not be performed mechanically. Data must be comparable in operating context, censoring must be handled correctly, and mixed failure modes can distort the fitted shape. A single beta value does not identify the physical mechanism. Engineering evidence is needed to connect the statistical pattern to manufacturing defects, random overstress, fatigue, corrosion, wear, or another cause.

If failures are pooled across several mechanisms, the apparent population may not fit a simple Weibull interpretation. Segregating by failure mode, equipment variant, duty, environment, or configuration can reveal more useful patterns. Standardized asset and event data are therefore prerequisites for credible reliability modeling.

WEIBULL SHAPE	QUALITATIVE INTERPRETATION
$\beta < 1$	Decreasing hazard; investigate early-life / latent-defect mechanisms.
$\beta = 1$	Constant hazard; fixed-age replacement needs a separate technical justification.
$\beta > 1$	Increasing hazard; investigate age-related degradation and the effectiveness of intervention.

2.6 Load, environment, derating, and margin

Reliability is strongly influenced by stress. Temperature, voltage, current, pressure, speed, vibration, contamination, humidity, corrosive exposure, cyclic load, software demand, and mechanical shock can change failure behavior. A component that is reliable at nominal laboratory conditions may perform differently in a hot enclosure, high-vibration vehicle, coastal environment, or frequently cycled process.

Derating reduces operating stress relative to rated capability. The appropriate margin depends on the component, failure mechanism, applicable standard, environment, and mission. Excessive margin can increase size and cost without proportional benefit; inadequate margin accelerates degradation or leaves little tolerance for uncertainty. The engineering objective is justified margin, not maximum conservatism.

Environmental qualification and duty characterization should therefore be part of reliability planning. If the actual duty is poorly known, operating data can be used to validate assumptions after commissioning. A reliability model based on one load spectrum should be revisited when service conditions change.

2.7 Hidden functions and protective systems

Some functions are hidden during normal operation. A relief valve, emergency shutdown, standby pump, backup power supply, fire damper, protection relay, or fail-safe mechanism may be failed without any immediate process symptom. The failure becomes evident only when the protective function is demanded or explicitly tested.

Hidden failures require a different maintenance logic. Waiting for functional failure may be unacceptable because the failure can remain latent until another event occurs. Failure-finding tasks are designed to reveal whether the protective function is available. The test interval should reflect consequence, demand frequency, failure likelihood, testing burden, and any governing requirement.

Testing itself can create risk. Overly frequent proof testing can cause wear, introduce configuration error, or expose the system to transition hazards. The task should be technically effective and proportionate, and the test should restore the system to the correct configuration afterward.

> **Worked example.** Two independent full-duty pumps each have mission reliability 0.90, giving ideal parallel reliability $1 - (1 - 0.90)^2 = 0.99$. If both require one motor-control center with mission reliability 0.95 and that support is independent of pump failures, complete-path mission reliability is $0.95 \times 0.99 = 0.9405$. The shared support bounds the result at 0.95; additional dependencies can reduce it further.

MODULE 3 | 25 MINUTES

Maintainability, Maintenance Strategy & Condition Monitoring

3.1 Maintainability begins in design

A maintainable design allows personnel to inspect, isolate, diagnose, remove, repair, replace, adjust, test, and return the asset to service safely and efficiently. Poor maintainability creates long outages, workarounds, error opportunities, and deferred maintenance even when the equipment itself is reliable. NASA guidance emphasizes maintainability as a design and lifecycle characteristic, not merely a repair statistic.

Access is one of the most visible maintainability features. Engineers should consider working clearance, reach, lighting, lifting paths, removable panels, rigging points, drain locations, isolation points, hazardous-energy control, environmental protection, and whether adjacent equipment must be disturbed. A filter that technically can be replaced but requires removal of piping and a temporary platform is maintainable only in a very limited sense.

Diagnostic design is equally important. Fault codes should discriminate among likely causes rather than report a generic “system fault.” Test points, trends, self-checks, logs, and clear alarm priorities can reduce troubleshooting time. Diagnostic coverage should also consider sensor failure; an incorrect indication can misdirect maintenance and increase downtime.

MAINTAINABILITY FACTOR	DESIGN / PROCUREMENT CHECK
Access / isolation	Safe clearances, lifting, isolation and work sequence.
Diagnostics / modularity	Locate the fault, replace the intended module, prevent connection errors.
Standardization / documentation	Compatible spares, current drawings, procedures and configuration control.

3.2 MTTR is a system property, not just wrench time

Mean time to repair (MTTR) is often used as a maintainability metric, but organizations define its boundaries differently. One metric may measure hands-on repair only. Another may include detection, dispatch, access, diagnosis, parts retrieval, permits, cooling time, repair, testing, and return to service. Comparing MTTR across assets or organizations is meaningless unless the time boundary is defined.

For availability decisions, total restoration time is usually more important than hands-on task duration. A failed component may require only thirty minutes to replace but remain unavailable for two days while a spare arrives. Conversely, a planned outage may require eight hours of labor but little service interruption because work is coordinated with scheduled downtime.

Engineers should decompose restoration time into actionable segments. Long diagnosis suggests better diagnostics or training. Long parts delay suggests a spare or supplier problem. Long access time

suggests design changes. Long administrative delay may point to permitting, approval, or planning. The purpose of the metric is to improve the system, not simply report an average.

3.3 Reactive maintenance has a legitimate place

Run-to-failure is appropriate for some assets. If failure has low safety, environmental, operational, and secondary-damage consequence; failure is evident; restoration is fast; spares are inexpensive; and no preventive task is technically effective, deliberate reactive maintenance may be the most economical strategy. The key word is deliberate.

Unplanned reactive culture is different. When critical assets are allowed to fail because maintenance was deferred, known degradation was ignored, or data were unavailable, the organization is accepting risk without a structured decision. High reactive workload also consumes labor that could be used for planning, condition monitoring, and defect elimination, creating a reinforcing cycle.

A reliability-centered program therefore does not seek zero reactive work. It seeks an optimized mix in which run-to-failure is applied intentionally to appropriate failure modes while consequential failures are managed through effective tasks, monitoring, redesign, redundancy, or other controls.

3.4 Time-based preventive maintenance must match age behavior

Preventive maintenance replaces, restores, lubricates, calibrates, or inspects an item at a time or usage interval. It is effective when the task addresses a failure mode with a meaningful relationship to age or usage, when the interval can be set before unacceptable probability or degradation develops, and when the task itself restores resistance to failure.

Some maintenance is mandatory regardless of a purely economic reliability analysis. Codes, regulations, manufacturer requirements, warranties, insurance conditions, or owner procedures may specify intervals. Where intervals are discretionary, operating evidence should be used to refine them. Performing a task too frequently wastes resources and can introduce infant mortality; performing it too late allows degradation to progress beyond a useful intervention point.

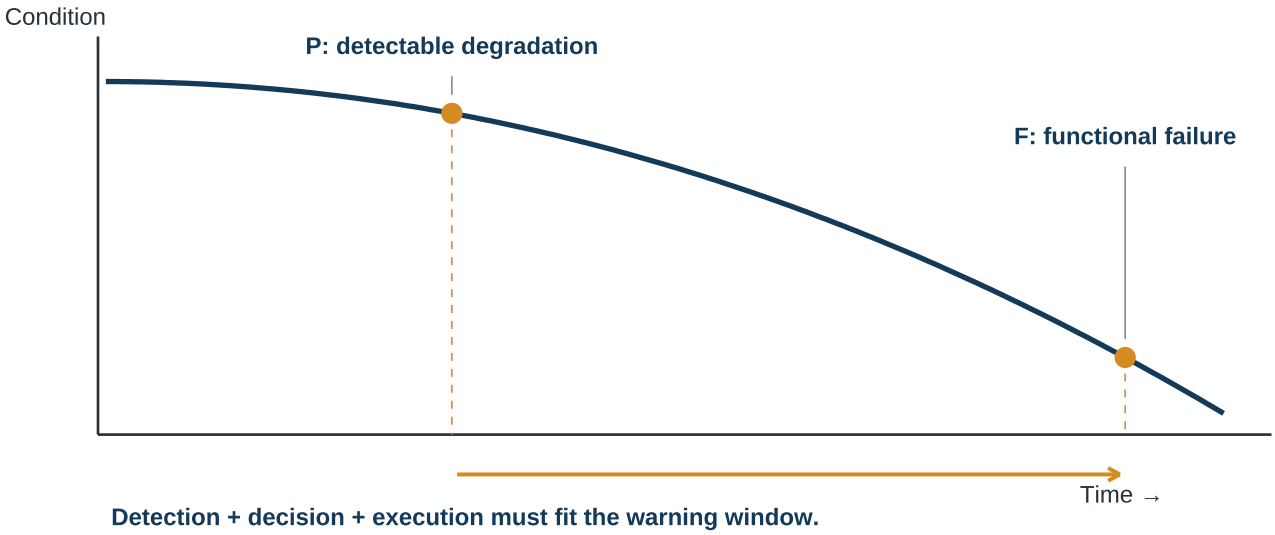
Age exploration can be used cautiously to refine intervals when data show that condition at the existing interval is consistently healthy and consequence permits extension. The process should have technical criteria, controlled changes, monitoring, and review rather than informal “we never find anything, so skip it” reasoning.

3.5 Condition-based maintenance uses the degradation signal

Condition-based maintenance acts on measured evidence of degradation rather than a fixed calendar alone. Techniques include vibration analysis, oil analysis, thermography, ultrasonic inspection, motor-current analysis, corrosion monitoring, performance trending, insulation testing, partial-discharge monitoring, thickness measurement, leak detection, and many discipline-specific methods.

The technology matters less than whether the parameter is connected to the failure mechanism and provides sufficient warning. A sensor that changes only moments before failure has little maintenance value unless automated protection can act. A parameter that drifts months in advance may support planned intervention. Data quality, sensor location, alarm thresholds, sampling rate, and environmental effects must be understood.

Condition monitoring can produce false alarms and false reassurance. Trend interpretation should consider normal operating changes such as load, speed, ambient temperature, process composition, and seasonal behavior. Baselines and normalization improve signal quality. Critical decisions should not rely on a single measurement without considering calibration, repeatability, and corroborating evidence.



A qualitative P-F interval connects detectable degradation to functional loss. The monitoring interval must leave time for diagnosis, planning and execution, with uncertainty considered.

3.6 The P-F interval guides inspection frequency

The potential-failure-to-functional-failure interval, commonly called the P-F interval, is the period between the point when degradation becomes detectable and the point when required function is lost. If an effective inspection method can reliably detect the condition during that interval, the inspection frequency can be selected so the organization has adequate probability of detection and enough time to plan and execute corrective action.

Inspection frequency should be meaningfully shorter than the P-F interval because detection is not guaranteed at the first possible instant. If the interval is highly variable, the task should be based on the shortest credible warning period or another conservative basis. The engineer should also consider inspection sensitivity, access, instrument capability, and whether the degradation can accelerate.

The P-F concept discourages arbitrary monitoring. Checking monthly because “monthly is common” is not technically justified unless it relates to degradation behavior. Conversely, a high-frequency sensor is not automatically useful if data are not reviewed or alarm thresholds do not drive action.

3.7 Maintenance-induced failures and post-maintenance testing

Maintenance can introduce failures through contamination, incorrect assembly, wiring errors, wrong software configuration, improper torque, omitted seals, damaged connectors, reversed valves, or incomplete restoration. The period immediately after intrusive maintenance can therefore have elevated risk.

Good maintainability design reduces the opportunity for error through keyed connectors, standardized parts, clear labeling, mistake-proofing, accessible inspection points, controlled procedures, and configuration management. Post-maintenance testing should verify the function affected by the work, not merely that power is present or the equipment starts.

When repeated failures follow maintenance, the reliability problem may be in the maintenance process rather than the underlying asset. Work-order history should distinguish failure discovered before maintenance from defects introduced during maintenance so that improvement efforts target the correct mechanism.

STRATEGY	WHEN IT MAY BE EFFECTIVE
Run to failure	Acceptable consequence with feasible restoration and no controlling requirement to intervene sooner.
Time / cycle based	Age-related mechanism and an intervention that restores resistance to failure.
Condition based	Detectable degradation with an actionable warning window.
Failure finding	Reveal a hidden functional failure before the combined consequence becomes unacceptable.
Redesign	Remove or reduce a mechanism or consequence that cannot be adequately managed by a task.

3.8 Planning and spare readiness are part of maintainability

A well-designed repair can still produce long downtime if the organization lacks the right spare, tooling, lifting equipment, procedure, trained personnel, or vendor access. Supportability therefore extends maintainability beyond physical design into the resources required to restore function.

Critical jobs should have realistic job plans that identify isolation, permits, parts, consumables, tooling, lifting, calibration, testing, and return-to-service criteria. The level of planning should match consequence and complexity. Emergency plans should be tested or walked down before failure occurs, especially for assets with long lead times or difficult access.

Spare strategy should consider failure probability, consequence, supplier lead time, repairability, shelf life, obsolescence, commonality, storage conditions, and whether a temporary operating configuration is available. The most expensive spare may be the one that is cheap to buy but impossible to obtain when the asset fails.

Applied evidence: from thermal image to a verified work record

The following published examples show why a thermal indication is a starting point for investigation. They do not establish a universal alarm temperature, a confirmed failure mechanism, or the condition of a BuilderCE project. The DOE guide presents paired infrared and visible-light views so the engineer can connect the thermal pattern to a physical component.

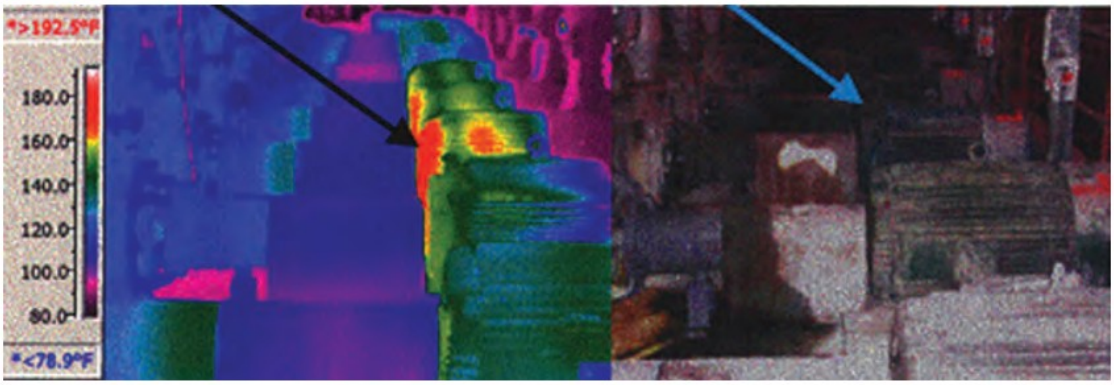


Figure 6.2.6. IR scans of multiple electric motors can highlight those with hot bearings indicting an imbalance or wear problem.

DOE O&M Best Practices Guide, Release 3.0 (August 2010), Figure 6.2.6, printed page 6.6. Original infrared / visible motor pair and source caption; reproduced for technical study.

Motor observation. The marked region is localized near the motor bearing area, while the visible view locates the component in the installation. DOE describes a possible bearing-related issue. Color alone cannot distinguish loading, lubrication, alignment, bearing wear, reflected radiation or other influences. Record the actual operating state and obtain corroborating evidence before assigning a cause.

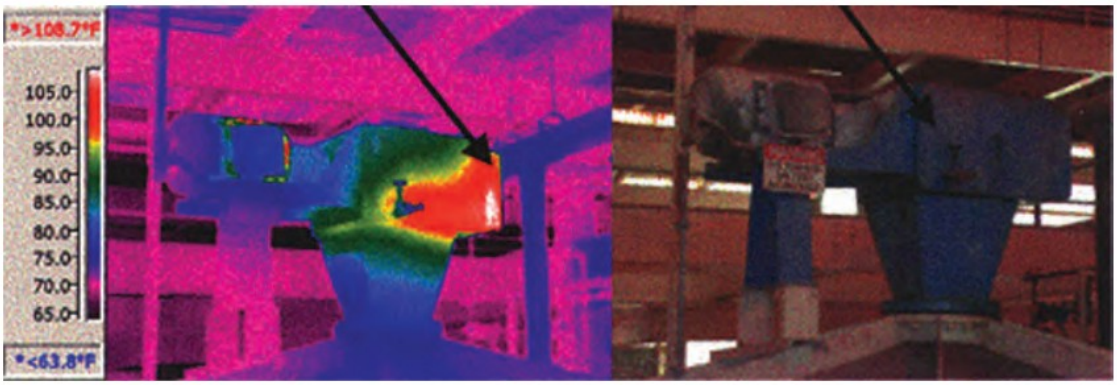


Figure 6.2.7. Possible gearbox problem indicated by white area defined by arrow. Design drawings of gearbox should be examined to define possible cause of elevated temperatures.

DOE O&M Best Practices Guide, Release 3.0 (August 2010), Figure 6.2.7, printed page 6.6. Original infrared / visible gearbox pair and source caption; DOE calls for drawing review to investigate the indicated region.

Gearbox observation. The arrow identifies a localized warm region, and the visible image identifies the corresponding geometry. Consult the gearbox drawing, bearing and lubrication arrangement before choosing the next check. The two published examples are different assets, not a before-and-after repair sequence. A color difference between them is not evidence of improvement.

RECORD BEFORE INTERPRETING	WHY IT MATTERS
Load, speed, duty and steady-state status	A hotter loaded asset cannot be directly compared with an unloaded or warming-up asset.
Ambient conditions, surface and measurement setup	Emissivity, reflected temperature, viewing angle, distance and camera settings affect the interpretation.
Comparable baseline and corroboration	Use repeatable locations and comparable duty; review vibration, lubrication, alignment and inspection evidence.
Consequence and response window	Use approved asset-specific criteria to decide monitoring, planned intervention or immediate operational escalation.

Completed diagnostic work request: WR-017

> **Illustrative training record.** All asset identifiers, actions and outcomes below are hypothetical. They are separate from the published DOE equipment. No plant temperature, test result or alarm threshold is inferred from those images.

RECORD FIELD	COMPLETED ENTRY
Function / context	P-201A process pump must deliver the approved duty at normal operating speed. P-201B is available as a verified standby; loss of both interrupts the process.
Observation / uncertainty	Repeated thermal indication at the drive-end bearing region at comparable load. Cause unconfirmed. No temperature is transcribed from the DOE illustrations.
Containment / owner	Operations confirms standby readiness and checks approved escalation criteria. Reliability engineer owns diagnosis; maintenance planner coordinates isolation and outage.
Diagnostic scope	Repeat thermography with the same setup; record actual load and ambient conditions. Review vibration spectrum, lubrication condition and history, coupling alignment and base / piping condition.
Evidence / diagnosis	In this example, the repeat scan and vibration pattern persist at comparable duty. An isolated alignment check identifies an out-of-tolerance coupling condition against the applicable manufacturer procedure. Lubrication checks do not support the initial over-greasing hypothesis.
Authorized action	Under the approved job plan, isolate the train, correct alignment and verify base condition. Inspect the bearing and coupling; renew components only as supported by the findings and manufacturer criteria. Preserve as-found / as-left records.

Verification, closure and learning

STAGE	COMPLETED EVIDENCE / DECISION
Before return to service	Maintenance records as-left alignment within the manufacturer limits and checks guards, isolation restoration and the agreed functional test. Operations authorizes the return to duty.
Comparable-duty recheck	The same observation points and measurement setup show that the localized anomaly is no longer present; corroborating vibration is within the approved asset criteria. Actual readings and attachments belong in the work record.
Recurrence check	After 30 operating days at the relevant duty, the repeat indication and vibration pattern have not recurred. This supports the local correction; it does not establish a new MTBF or prove indefinite reliability.
Record / reopen trigger	WR-017 links the scans, vibration report, alignment record, work order and operator release. Reopen on renewed indication, changed duty, abnormal vibration or a relevant configuration change.

This record separates observation, hypothesis, confirmed finding, action and verification. A work order closed after parts replacement alone would not establish that the recurring mechanism had been controlled. The comparison is meaningful because the duty and measurement conditions are preserved, and the verification checks the same functional concern that justified the work.

Primary evidence:DOE O&M Best Practices Guide, Figures 6.2.6–6.2.7

MODULE 4 | 25 MINUTES

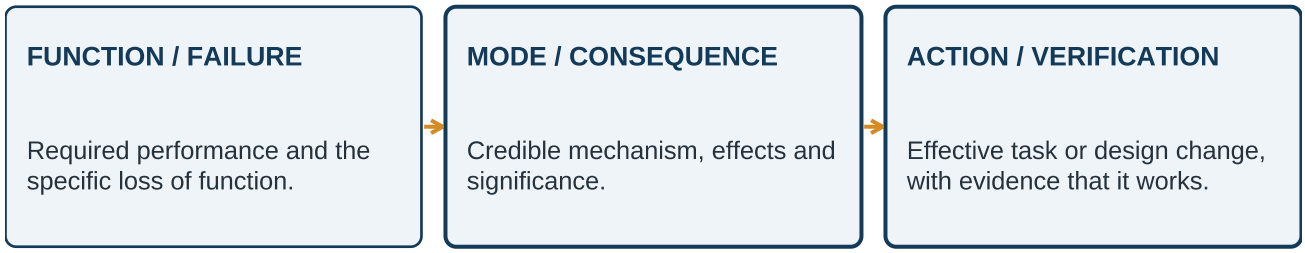
Reliability-Centered Maintenance, Criticality & Asset Strategy

4.1 RCM begins with operating context and function

Reliability-centered maintenance is a structured process for determining how failure should be managed so an asset continues to fulfill required functions in its operating context. Current DOE guidance describes RCM as a strategic combination of maintenance approaches coupled with root-cause analytics, while NASA applies RCM broadly to facilities and collateral equipment. SAE JA1011 defines evaluation criteria for processes that claim to be RCM.

RCM should not begin with the question “What maintenance does this equipment need?” It begins with “What functions must this asset perform, to what performance standard, and in what context?” The analysis then identifies functional failures, failure modes, effects, consequences, and technically effective actions. This sequence helps prevent maintenance tasks that are based on tradition rather than failure behavior.

Operating context can change the answer. A small pump may be run to failure when a redundant pump is available and the consequence is minor, but the same pump model may require condition monitoring and stocked spares when it supports a single critical process. RCM is therefore function- and context-specific, not a generic maintenance schedule for an equipment class.



RCM connects required function, functional failure, failure mode and consequence to a technically effective failure-management action.

A useful failure mode describes the mechanism or condition that causes functional loss at a level that supports a maintenance or design decision. “Pump failure” is too broad. “Bearing fatigue,” “seal damage from dry running,” “motor insulation breakdown,” “impeller erosion,” and “loss of suction from clogged strainer” suggest different detection methods and tasks.

Too much detail is also inefficient. An analysis that lists every microscopic mechanism can become unmanageable without improving decisions. The appropriate level is the level at which the failure-management strategy changes. If several causes are detected and managed by the same task, grouping may be reasonable. If one cause requires redesign and another is adequately managed by monitoring, they should remain distinct.

Failure modes should include credible human and process causes where relevant, such as incorrect configuration, contamination during maintenance, software update error, or loss of calibration. These are engineering reliability issues when they can defeat required function.

4.3 Consequence determines how much reliability effort is justified

Not all failures deserve the same analysis depth. Criticality combines the significance of consequence with the likelihood or frequency of failure and often includes detectability, recovery, or other factors. The purpose is prioritization, not creation of a mathematically exact risk number.

Consequence categories may include safety, environmental impact, mission or service loss, secondary equipment damage, quality, regulatory impact, and cost. A low-cost component can be highly critical if its failure disables a critical system or causes long downtime. Conversely, an expensive component may be suitable for run-to-failure when redundancy and rapid replacement make the operational consequence small.

Criticality should also influence data quality expectations, spare strategy, monitoring depth, review independence, testing, and change control. High-consequence systems justify stronger evidence and more conservative assumptions than low-consequence convenience equipment.

CRITICALITY CONSIDERATION	DECISION SUPPORTED
Safety / environment / service	Depth of analysis, review independence and treatment priority.
Likelihood / detection / restoration	Monitoring, hidden-function tests, spares and response capability.
Uncertainty / changed duty	Conservative assumptions, targeted evidence and a trigger to revalidate.

4.4 Hidden failures and multiple-failure consequence

A hidden protective failure may have little consequence by itself but creates exposure to a second event. For example, a failed standby pump is harmless while the primary pump operates, but consequential when the primary fails. A protection relay that is failed but undetected becomes important when the protected fault occurs. RCM therefore evaluates multiple-failure consequence for hidden functions.

Failure-finding intervals should be based on the reliability of the hidden function, demand frequency, consequence of combined failure, and test effectiveness. A test that does not actually exercise the critical success path may provide false assurance. For example, starting a generator without transferring load does not verify all functions required for emergency service.

The analysis should also consider whether the test can reveal dormant common-cause problems. A common fuel or control dependency may remain untested if each redundant unit is exercised in isolation under ideal conditions.

4.5 Spares are reliability controls, but not free controls

Spares reduce restoration time when failure cannot be fully prevented or predicted. Their value rises with long procurement lead time, high downtime consequence, low interchangeability, and limited repair alternatives. But inventory creates capital cost, preservation requirements, shelf-life risk, obsolescence, configuration-control burden, and the possibility that the stored item is unavailable or unusable when needed.

Engineers should distinguish consumables, insurance spares, rotable repairables, strategic modules, and commodity items. A critical spare may justify environmental control, periodic inspection, firmware preservation, battery maintenance, or documented interchangeability. A spare that degrades in storage is not a dependable mitigation.

Commonality can reduce spare burden but may increase common-cause vulnerability if many systems depend on one model, supplier, software version, or design defect. Standardization and diversity should therefore be balanced rather than treated as universally good or bad.

SPARE STRATEGY	CHECK BEFORE CREDITING THE CONTROL
Exposure / lead time	Failure consequence, supplier uncertainty and temporary operating options.
Preservation / compatibility	Shelf life, storage condition, accessories, firmware and interchangeability.
Logistics / restoration	Transport, lifting, tools, trained personnel and a realistic installation plan.

4.6 Repair-versus-redesign decisions

Repeated corrective maintenance is often a symptom that the failure mode is not being managed effectively. If no preventive or condition-based task can reduce the likelihood or consequence to an acceptable level, redesign may be the appropriate reliability action. Redesign can remove a failure mechanism, reduce stress, add independence, improve diagnostics, provide access, or make restoration faster.

Redesign decisions should consider lifecycle exposure rather than only the next repair cost. A recurring failure that causes minor repair cost but major service disruption may justify capital improvement. Conversely, a redesign that is technically elegant but eliminates only low-consequence nuisance failures may not be economically justified.

Root-cause analysis is most valuable when it identifies a controllable mechanism and leads to verified corrective action. Merely replacing failed parts without addressing alignment, contamination, transient loading, environmental exposure, software logic, or another cause can create an endless repair cycle.

FAILURE-MANAGEMENT ACTION	TECHNICAL PURPOSE
Eliminate / reduce the mechanism	Change the design or operating condition that creates the failure.
Detect / intervene	Find actionable degradation and complete effective work in time.
Mitigate / restore	Limit consequence and restore required function using a justified support plan.

4.7 Performance contracts and procurement requirements

Reliability and maintainability can be specified during procurement, but requirements must be measurable and connected to operating context. Statements such as “high reliability” or “easy to maintain” are not enforceable. Better requirements define mission duration, demand profile, required availability, maintainability demonstration, maximum restoration time, diagnostic capability, access, spare support, environmental conditions, and data deliverables where appropriate.

Performance requirements should avoid forcing a vendor to produce a single MTBF number that cannot be validated. Reliability evidence may include analysis, historical data from comparable service, qualification testing, design margins, FMEA, accelerated testing, or reliability demonstration. The required evidence should match consequence, novelty, and the maturity of the technology.

Maintainability demonstrations can reveal design problems before commissioning. Representative personnel can perform defined tasks using intended tools, procedures, and access. Observed task time, error opportunities, access constraints, and support needs can then be corrected while design changes remain practical.

4.8 Asset strategy should be periodically revalidated

Maintenance strategy is not permanent. Duty changes, new failure mechanisms, aging, supplier changes, obsolescence, new monitoring capability, changes in consequence, and actual operating data can make the original strategy obsolete. A pump moved from intermittent standby to continuous service may require a different approach. A component with repeated infant mortality after overhaul may justify changing the overhaul process rather than the interval.

Periodic review should focus on material changes and bad actors rather than mechanically reanalyzing every asset. Trigger-based review can be initiated by recurring failures, unacceptable availability, missed condition detections, new hazards, major design changes, long-lead obsolescence, or evidence that a preventive task is not finding meaningful degradation.

The objective is a living failure-management strategy tied to evidence. RCM is strongest when it becomes part of engineering decision-making rather than a one-time workshop document stored after implementation.

> **Practice principle.** A maintenance task is justified because it is technically effective for a specific failure mode and worthwhile for the consequence - not because it is preventive, predictive, or traditional.

MODULE 5 | 20 MINUTES

Reliability Data, Growth & Applied Engineering Decisions

5.1 Reliability data must preserve context

Reliability analysis is only as credible as the data. ISO 14224 provides an industry-specific example of structured reliability and maintenance data, emphasizing equipment taxonomy, failure data, maintenance data, downtime, and data quality. Even outside petroleum and process industries, the

underlying lesson is broadly applicable: asset, event, and maintenance records must use consistent definitions if they are expected to support analysis.

A failure record should identify the asset configuration, operating context, symptom, functional impact, failure mode where known, cause where supported, consequence, dates and times, maintenance action, parts used, labor, and return-to-service status. Not every organization needs every field, but the data should support the decisions being made.

Free-text work orders contain valuable detail but are difficult to trend. Coded fields support analysis but can oversimplify. A strong system uses controlled codes for key attributes and retains narrative evidence for unusual or complex events. Data governance should define who can change failure classifications and how corrections are preserved.

DATA FAMILY	CONTEXT TO PRESERVE
Asset / configuration	Identity, boundaries, duty, version and operating environment.
Event / cause	Time, symptoms, functional effect, supported mode and confirmed cause where known.
Maintenance / verification	Action, resources, parts, restoration time, as-left checks and recurrence.

5.2 Bad averages hide reliability problems

MTBF can improve while a subset of critical assets deteriorates. Availability can remain high because redundant capacity masks repeated failures. Preventive-maintenance compliance can be 98 percent while the two percent missed consists of the most consequential work. Dashboard metrics therefore need engineering interpretation.

Distributions and segmentation are often more useful than averages. Analyze by asset family, failure mode, duty, site, manufacturer, age, environment, maintenance team, or configuration. Separate planned from unplanned downtime. Track repeat failures and infant mortality after maintenance. Examine high-duration events even when they are infrequent.

Metric definitions should remain stable over time. Changing the clock boundary for MTTR or reclassifying failures can create an apparent improvement without a physical change. If definitions change, the historical series should be reconciled or clearly annotated.

MEASURE	INTERPRET WITH
MTBF / MTTR / availability	Consistent clocks, comparable duty, distributions and uncertainty.
PM compliance / condition alerts	Consequence of missed work, signal reliability and completed action.
Repeat failures	Asset boundary, mode and recurrence window; investigate incomplete defect elimination.

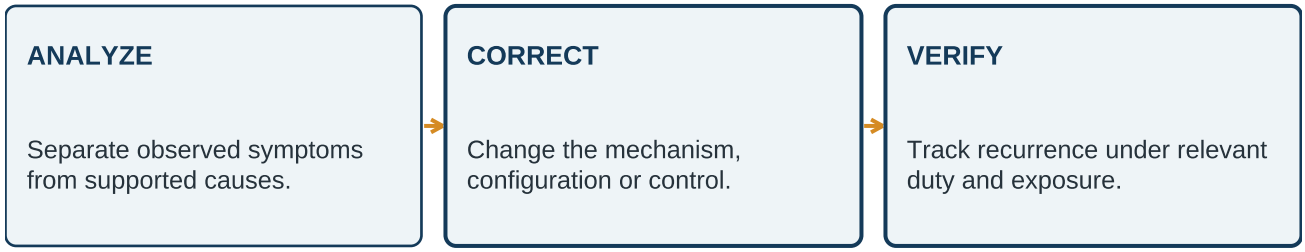
5.3 Reliability growth requires closed-loop corrective action

Reliability growth occurs when failures are discovered, analyzed, corrected, and the improvement is verified. Testing alone does not create growth. Repeatedly exposing the same design to failure without removing causes only accumulates evidence of weakness.

Growth programs identify failure modes, prioritize them by consequence and recurrence, implement design or process changes, and then evaluate whether the failure rate or mechanism has changed. Corrective action can involve design margin, material, supplier process, software, installation,

maintenance, operating procedure, training, or inspection. The corrective action should address cause, not merely symptom.

Verification should look for recurrence under relevant conditions. A single successful restart after a repair does not demonstrate long-term reliability improvement. Trend monitoring, additional testing, fleet comparison, or a defined period of successful operation may be needed depending on consequence and exposure.



Reliability growth requires a change to the failure process and verification under relevant exposure.

5.4 Physics of failure and statistical evidence are complementary

Statistical reliability methods describe patterns in observed or modeled failure data. Physics-of-failure methods examine the actual degradation and failure mechanisms: fatigue, creep, corrosion, wear, thermal cycling, electromigration, insulation breakdown, chemical attack, vibration, or other processes. NASA maintains a current Physics of Failure Handbook as part of its reliability and maintainability technical program.

Physics-based reasoning is particularly useful when historical data are sparse, the design is new, or operating conditions differ from the population that produced the data. Statistical evidence is valuable for detecting fleet patterns and quantifying experience. The strongest reliability decisions use both when appropriate: engineering mechanism explains why, while data helps show how often and under what conditions.

Neither approach should create false certainty. A detailed model still depends on material properties, loads, boundary conditions, manufacturing variability, and model validity. Field data can be biased by incomplete reporting or changing configurations. Uncertainty should remain visible in the decision.

5.5 Applied case - recurring process-pump failures

A process pump experiences three bearing and seal failures in eighteen months. Maintenance replaces the bearing and seal each time. Vibration data show rising axial vibration before the latest event, and operators report intermittent low suction pressure during high-demand periods. The repeated “bearing failure” code has led management to consider stocking more bearings.

A reliability-focused review begins with function and operating context. The pump must deliver a defined flow and head across a variable process demand. The team reviews alignment records, base condition, piping strain, suction pressure, net positive suction head margin, lubrication, vibration spectrum, seal flush conditions, operating point relative to the pump curve, and maintenance practices. The evidence indicates that the pump frequently operates far from its preferred range and experiences cavitation during demand transients.

The most effective action is not a larger spare inventory. The team corrects the suction restriction, modifies operating logic to reduce off-design operation, verifies alignment, establishes condition thresholds, and retains an appropriate spare kit. The reliability program then tracks vibration, seal leakage, and recurrence. This closes the loop from symptom to mechanism to verified improvement.

5.6 Applied case - emergency generator availability

A facility has two emergency generators, each capable of carrying the critical load. Historical engine reliability is strong, but several events involve failure of starting batteries and one involves a common

fuel-quality problem. Restoration following battery failures was rapid; the fuel event disabled both generators until temporary fuel treatment was arranged.

A simple MTBF report can hide the most important reliability issue. The common fuel dependency dominates the multiple-failure consequence even though engine failure frequency is low. The engineering response may include fuel-quality monitoring, tank maintenance, separation or diversity of supply, periodic loaded testing, battery condition monitoring, and verification of automatic transfer functions. The selected controls should address the complete emergency-power success path.

Availability planning should also include restoration. If the facility cannot tolerate extended loss, a connection for a mobile generator, documented hookup procedure, tested switching arrangement, and vendor response agreement can reduce consequence. Recovery capability is part of reliability engineering when it determines whether required service can be restored within the mission need.

5.7 Applied case - aging movable bridge machinery

A movable bridge has mechanical drive components, brakes, limit switches, controls, position feedback, and standby power. Several parts are obsolete, but major rehabilitation is planned three years in the future. The owner must decide which components to replace now, which to monitor, and which to carry as strategic spares.

Criticality analysis identifies failure modes that can strand the bridge, create unsafe movement, or require long closure. Condition inspections identify gear wear and brake degradation with measurable trends. Some electronic modules have little warning before failure and long procurement lead times. The strategy therefore mixes condition-based tasks for degrading mechanical elements, functional testing for hidden protective devices, strategic spares for obsolete electronics, and defined emergency operating procedures.

The planned rehabilitation changes the economic horizon but does not eliminate reliability obligations during the remaining service period. Temporary strategies should be documented with their assumptions, triggers, and expiration so they do not silently become permanent.

5.8 Applied case - building automation and common software dependency

A large building uses redundant servers for building automation, but both run the same software version, depend on a common database, and are administered through the same identity service. Mechanical equipment can operate locally, but central scheduling, trending, alarms, and optimization depend on the shared digital architecture.

The system has hardware redundancy but significant common software and service dependencies. Reliability assessment should identify which functions truly require central control and which can fail gracefully to local operation. Backup configuration, tested restore procedures, network segmentation, local fallback, spare controllers, and change control may contribute more to availability than another identical server.

This case illustrates the breadth of reliability engineering. Failure mechanisms can be mechanical, electrical, software, human, environmental, or organizational. The required function and recovery path determine which dependencies matter.

5.9 Applied case - spare transformer strategy

A utility has several large transformers with long procurement lead times. Failure probability is low, but consequence includes prolonged outage and major replacement logistics. Purchasing a full spare for every location would be expensive and may create storage and preservation problems.

The engineering analysis considers fleet commonality, loading, condition, age, transport routes, bushings and accessories, failure history, supplier lead time, repair capability, temporary configurations, and whether one strategic spare can serve multiple sites. Standardizing interfaces can improve interchangeability, while transport and installation studies verify that the spare is actually deployable.

The decision is not simply probability times replacement cost. Long-duration service loss, restoration logistics, regulatory or customer consequence, and the option value of a shared strategic spare can dominate. Reliability and lifecycle economics should therefore be connected for major asset decisions.

5.10 Final reliability and maintainability checklist

- Required functions and performance standards are explicit for the actual operating context.
- Failure data distinguish symptom, failure mode, cause, consequence, and exposure basis.
- MTBF/MTTF, MTTR, availability, and failure-rate metrics use consistent definitions and time boundaries.
- Redundant paths have been checked for shared power, environment, software, maintenance, and other common causes.
- Maintainability has been addressed through access, isolation, diagnostics, modularity, tooling, documentation, and support.
- Preventive tasks have a credible age/usage basis or governing requirement; condition-based tasks have a useful detection interval.
- Hidden protective functions have appropriate failure-finding tests and restoration checks.
- Criticality drives analysis depth, spare strategy, monitoring, review, and escalation.
- Recurring failures receive mechanism-focused corrective action rather than repeated part replacement alone.
- Operating and maintenance data feed back into design, procurement, maintenance intervals, and asset strategy.

> **Practice principle.** Reliability and maintainability are not maintenance statistics. They are engineering characteristics created by function, architecture, design margin, dependencies, access, diagnostics, support, maintenance strategy, and continuous learning from operating evidence.

Advanced Practice - Reliability Decisions That Commonly Go Wrong

A.1 Treating MTBF as guaranteed life

An MTBF is a population average under a defined model and exposure. It is not a promised failure-free life or the age at which each individual component will fail. Use a reliability function or age-dependent model when mission survival is the real question, and state the operating conditions that support the data.

A.2 Counting redundant components instead of independent success paths

Two pumps, two servers, or two feeders do not provide two independent success paths when they share suction, controls, power, environment, software, or maintenance. Reliability diagrams should include the support functions that can defeat the redundancy.

A.3 Scheduling preventive replacement without age evidence

Calendar replacement is not automatically preventive in the technical sense. If the failure mode is random or dominated by installation error, periodic replacement can create cost and new defects without meaningfully reducing failure probability.

A.4 Monitoring a parameter that does not provide actionable warning

Condition monitoring must be tied to a failure mechanism and a P-F interval long enough to act. Collecting high-frequency data that nobody reviews or that changes only seconds before failure creates instrumentation, not maintenance strategy.

A.5 Improving repair time while ignoring recurrence

Faster restoration improves availability, but repeated failure may still be unacceptable. A strong program separates restore-service work from defect-elimination work so the same failure does not become a permanent workload.

A.6 Confusing maintenance compliance with reliability

Completing every scheduled task does not prove the tasks are effective. Examine whether the tasks find meaningful degradation, prevent or detect the targeted failure modes, and improve performance without creating maintenance-induced failures.

A.7 Using poor-quality data to create precise statistics

A failure database with inconsistent equipment boundaries, failure definitions, and downtime clocks can generate precise but misleading MTBF and availability values. Data-definition discipline is an engineering control.

A.8 Deferring obsolescence until failure

Obsolescence can turn a minor failure into a long outage. Strategic reliability planning identifies unsupported hardware, software, proprietary interfaces, long lead times, and disappearing repair capability before the failure occurs.

A.9 Letting the maintenance strategy outlive the operating context

Duty, environment, redundancy, consequence, and technology change. A task selected for intermittent duty may be wrong after the asset becomes continuous service. Strategy should be revalidated when context changes.

A.10 Declaring success without verifying corrective action

A redesign or maintenance change should be followed by evidence that recurrence changed. Reliability improvement is not complete when the work order closes; it is complete when the targeted failure process is demonstrably controlled.

Technical Deep Dives - Reliability & Maintainability Practice

B.1 Mission reliability versus long-term availability

Mission reliability and availability answer different questions and can lead to different architecture choices. Mission reliability asks whether the system can perform for a specific uninterrupted period without failure. Availability asks whether the system is ready for use when required over a longer operating horizon that can include failures and restoration. A system that fails relatively often but is repaired quickly can achieve high availability while having poor mission reliability. Conversely, a system that almost never fails but requires weeks to repair can have high mission reliability over short missions but poor long-term availability after a rare event.

The distinction matters in standby and emergency systems. An emergency generator may need very high probability of successful start and sustained operation for a defined outage duration. Repair after the outage does not help the mission that already failed. For a continuously operated production system with redundant trains, short repair times and switching capability may be more important because service can continue while one train is restored.

Engineers should therefore state the decision metric explicitly. If the requirement is “operate continuously for 72 hours after loss of normal power,” use a mission-reliability framework for the complete emergency-power path. If the requirement is “provide at least 99.9 percent service over the

year," availability and maintainability become central. The same asset can have both requirements, and one metric should not be used as a substitute for the other.

B.2 Repairable-system failure processes

Repairable assets produce sequences of failures and restorations rather than one terminal life. The failure process may be stable, improving, or deteriorating. A simple average MTBF assumes that the underlying process is sufficiently stationary for the average to be meaningful. When modifications, aging, duty changes, or recurring defects alter the process, the historical mean can hide the trend.

For a repairable system, engineers should examine failure times in sequence. Are intervals between failures getting shorter? Did a design modification change the pattern? Are several events repeats of one failure mode? Does the asset experience bursts of failures after maintenance? Trend plots and cumulative failure counts can reveal behavior that an overall MTBF conceals.

A recurring-repair cycle should trigger a different question from an isolated random event. The organization may need root-cause analysis, redesign, changed operating limits, improved installation, or a different maintenance task. Reliability growth methods are particularly relevant when corrective actions are intentionally introduced and subsequent operation is used to verify improvement.

B.3 Confidence and uncertainty in reliability estimates

Reliability statistics calculated from limited experience contain uncertainty. Ten assets operating for a few months provide less evidence than a mature fleet with years of exposure, even if both produce the same point estimate. Zero failures does not prove perfect reliability. It may simply mean the accumulated exposure is too small to observe a rare event.

When a decision is consequential, the engineer should consider confidence bounds, conservative estimates, or scenario ranges rather than present one point value as certain. The specific statistical method depends on the assumed failure process, censoring, sample size, and available data. The principle is broader: confidence should reflect the evidence.

Sparse data can be supplemented with qualified analogous experience, manufacturer evidence, accelerated testing, physics-of-failure analysis, and expert judgment, but the transferability of each source should be explained. A component used at lower temperature and duty in another industry may not provide a directly applicable failure rate. Combining several weak sources does not automatically produce strong evidence.

B.4 Censoring and incomplete life data

Reliability datasets often include assets that have not failed by the end of the observation period. Those assets are right-censored: they provide evidence that life exceeds the observed duration but do not provide a failure time. Discarding them biases life estimates downward or otherwise distorts the analysis. Conversely, recording only failures without the surviving population prevents a valid exposure-based estimate.

Maintenance databases are rarely organized as perfect statistical experiments. Assets enter and leave service, configurations change, and failure definitions evolve. Before fitting a distribution, the engineer should understand which units were at risk, for how long, and whether preventive replacement removed units before failure. The data process can matter as much as the mathematical method.

Where the available data cannot support a formal life distribution, a simpler engineering summary may be more defensible. The analysis can identify observed failure ages, surviving exposure, dominant mechanisms, and sensitivity of the decision to plausible life ranges without claiming statistical certainty that the records do not support.

B.5 Demand-based protective reliability

Some protective components are not continuously exposed to failure in the same way as rotating equipment. They are demanded intermittently. A safety valve, emergency shutdown function, trip relay,

fire damper, standby pump, or backup system can remain latent until demanded. Reliability may therefore be expressed as probability of failure on demand rather than failures per operating hour.

Demand-based systems need proof testing or other evidence that hidden failures are found before a real demand. The proof-test interval, coverage, and restoration quality affect the probability that the protection is unavailable when needed. A test that exercises only part of the success path can leave hidden dependencies undetected.

Engineering judgment should consider demand frequency, failure mechanisms while dormant, degradation between tests, common-cause dependencies, and test-induced risk. A device that is mechanically exercised every year may still depend on an untested power supply or control signal. The protective function should be verified as a functional chain rather than a collection of independently tested components.

B.6 Reliability allocation and design requirements

System-level reliability objectives sometimes need to be allocated to subsystems or components so design teams can manage them. Allocation should reflect architecture, technology maturity, consequence, achievable performance, and the number of required elements. Equal allocation is easy but may be unrealistic if some elements are inherently more reliable or more difficult to improve.

Reliability requirements should include mission time, operating environment, duty, confidence expectations where appropriate, and the failure definition. “MTBF greater than 50,000 hours” without a stated basis can encourage paper compliance rather than useful engineering. The supplier may quote a prediction based on generic component handbooks that does not represent the actual operating environment or complete system.

Verification planning should be developed with the requirement. Analysis, similarity, test, demonstration, inspection, and operating evidence each have strengths and limitations. Very high reliability can require enormous test time to demonstrate statistically through failure-free testing alone, so a combined evidence approach may be necessary.

B.7 Maintainability allocation and task-time distributions

Maintainability requirements can also be allocated and demonstrated. A system may require that a defined percentage of corrective tasks be completed within a specified time using stated personnel, tools, procedures, and support. This is more informative than a generic statement that maintenance should be “easy.”

Repair times are distributions. A mean can be dominated by a few long events or hide that most tasks are fast but a small subset creates extreme outages. Median, percentile, and task-specific restoration times may therefore be more useful. Engineers should identify which tasks control service recovery for critical failure modes.

Demonstrations should use realistic constraints. If a maintenance trial is performed with unlimited access, expert designers, pre-positioned tools, and an open equipment panel, the result may not represent field maintainability. Representative personnel should perform representative tasks in the intended configuration whenever practical.

B.8 Graceful degradation and fail-safe behavior

Reliability engineering does not always require binary success or failure. Some systems can degrade gracefully, preserving a reduced but acceptable function after component failure. Examples include reduced pumping capacity, lower process throughput, restricted bridge operation, limited HVAC zones, degraded communication bandwidth, or a control system operating in manual mode.

Graceful degradation can improve functional availability when the reduced mode is intentionally designed, detectable, and operationally manageable. The degraded state should have defined limits, alarms, procedures, and restoration priorities. An unrecognized degraded state is not resilience; it is hidden failure.

Fail-safe behavior is different. A fail-safe design moves the system to a condition intended to minimize consequence when a failure occurs. The safe state may still represent loss of mission or service. Reliability and safety therefore interact but are not identical. Engineers should define what the system does when each important failure mode occurs and whether that response preserves safety, function, or both.

B.9 Reliability and human performance

Human actions influence reliability through operation, maintenance, calibration, configuration, diagnosis, and recovery. Some systems rely on an operator to recognize an alarm, select a backup path, isolate equipment, or execute a manual procedure within a limited time. That human action is part of the functional success path.

The engineer should evaluate whether the required action is detectable, understandable, feasible, practiced, and supported under the expected conditions. A procedure that works in daylight with full staffing may fail during an emergency with noise, protective equipment, or simultaneous alarms. Clear interfaces, alarm prioritization, labeling, physical differentiation, and automation can reduce dependence on memory and perfect execution.

Maintenance design should similarly reduce error opportunity. Keyed connectors, captive fasteners, standardized torque information, component identification, configuration checks, and post-maintenance functional testing are reliability controls because they reduce the probability that maintenance introduces a new defect.

B.10 Change control as a reliability control

Reliability assumptions can be invalidated by apparently minor changes. A software patch can alter timing. A replacement bearing can have different internal clearance. A vendor substitution can change material. A revised filter can alter pressure drop. An energy-saving control sequence can increase equipment cycling. A new network connection can introduce a common dependency.

Change review should therefore ask whether the modification changes stress, duty, interfaces, environment, failure modes, diagnostic coverage, spare interchangeability, or maintenance procedures. High-criticality assets may require formal reliability or FMEA review before implementation.

Post-change monitoring is important when the reliability effect is uncertain. Early-life failure trends, alarms, condition data, and maintenance burden can be compared with the pre-change baseline. The goal is not bureaucratic control of every change; it is preservation of the assumptions on which dependable performance relies.

Extended Applied Engineering Cases

C.1 Wastewater blower train

A wastewater treatment plant operates three large blowers, two required for peak aeration and one normally standby. Operators report that the standby unit often trips during monthly exercise because its inlet guide-vane actuator sticks. The trip is easily reset, so the event is repeatedly coded as a nuisance alarm rather than a failure.

The required function of the standby blower is not simply to rotate during a test. It is to start, synchronize with the process controls, and deliver required air when another blower is unavailable. The repeated actuator trips therefore represent functional unreliability of a redundant path. Because the plant requires two of three blowers during peak demand, the latent degradation also reduces system margin.

The reliability response includes actuator inspection, environmental review, lubrication or material assessment, exercise under representative load, and confirmation that the alarm and trip logic distinguish actuator position failure from other causes. A strategically stocked actuator may reduce restoration time, but eliminating the sticking mechanism is preferable if feasible. The maintenance

program should trend exercise failures separately from successful starts so the standby-path reliability is visible.

C.2 Traffic signal cabinet power supply

A corridor experiences intermittent traffic-signal dark conditions traced to aging cabinet power supplies. The supplies are inexpensive and can be replaced quickly, but failures create traffic disruption and field exposure for maintenance crews. The agency is considering periodic replacement of all units every five years.

A technically sound decision requires evidence that failure likelihood increases meaningfully with age or usage. The team segments the history by model, installation year, cabinet temperature, and site voltage conditions. The data show that failures are concentrated in one model exposed to high internal cabinet temperature rather than uniformly increasing after five years.

The improved strategy replaces the weak model at high-temperature sites, improves cabinet ventilation, stocks standardized spares, and uses periodic inspection for overheating indicators. A blanket five-year replacement program would have spent resources replacing many healthy units without addressing the dominant stress. Reliability analysis changed the maintenance decision by identifying mechanism and population rather than relying on age alone.

C.3 Chilled-water plant valve accessibility

A chilled-water plant uses several large isolation valves located above piping with minimal access. The valves are reliable, but when actuator or packing work is required, maintenance must build temporary scaffolding and remove nearby instrumentation. Restoration can take an entire shift, increasing outage risk during summer peak conditions.

The reliability problem is largely maintainability. During a planned plant upgrade, engineers add permanent access, lifting points, reposition selected instruments, and standardize actuator interfaces. The capital cost does not reduce the inherent failure rate of the valves, but it reduces restoration time, makes preventive inspection easier, and lowers the likelihood that maintenance is deferred because access is difficult.

Availability improves even if MTBF is unchanged. This case demonstrates why reliability programs that focus only on failure frequency miss important design opportunities.

C.4 Stormwater pump station and flood exposure

A stormwater pump station has two electrically driven pumps and one standby generator. The pumps are redundant, but the motor-control equipment and generator transfer switch are all located at the same low elevation. A severe flood could therefore disable both normal and emergency electrical paths.

A component-based reliability calculation may show high pump and generator reliability while missing the site-level common cause. The engineering assessment maps the complete success path and identifies flood elevation as a shared environmental dependency. Alternatives include elevating electrical equipment, adding flood barriers, relocating the generator and switchgear, or providing a portable connection at higher elevation.

The best option depends on flood hazard, consequence, access, lifecycle cost, and restoration planning, but the reliability principle is clear: redundancy must be protected from the hazards expected to create demand for it.

C.5 Laboratory exhaust fan maintenance

A research building has redundant exhaust fans serving critical laboratory spaces. One fan can support minimum safe flow while the other is maintained. Belt inspections are scheduled quarterly because that interval was inherited from an older facility. Technicians rarely find significant belt deterioration, and intrusive inspection itself requires temporary fan switching.

The team reviews belt failure history, operating hours, load, manufacturer guidance, condition indicators, and the consequence of an unexpected belt failure. They determine that a visual and vibration-based condition check can identify degradation without opening the guard at every quarter. Intrusive belt inspection is extended while loaded redundancy tests and differential-pressure monitoring continue.

The revised program reduces unnecessary intrusive work while preserving functional assurance. The change is documented as an age-exploration decision with defined indicators that would trigger a return to the shorter interval.

C.6 Bridge bearing replacement planning

Bridge inspection identifies several bearings with corrosion and restricted movement. The structure remains serviceable, but future movement demand and deterioration create uncertainty. Bearing replacement requires lane closures and jacking, so the owner must prioritize locations.

Reliability thinking begins with the function of each bearing and the consequences of restricted movement, not simply a corrosion rating. The team reviews movement demand, structural redundancy, inspection trend, temperature range, drainage, accessibility, and whether failure is detectable before structural consequence becomes significant. Bearings with rapidly worsening restriction and high structural consequence receive earlier intervention.

For lower-criticality locations, the owner may use condition monitoring and scheduled replacement during a future rehabilitation. The inspection interval is tied to observed deterioration and the warning period, demonstrating the P-F concept in a civil-infrastructure context.

C.7 Data-center UPS battery system

A data center has redundant uninterruptible power supply strings, each using many battery modules. Individual battery failures are relatively frequent, but the system remains available because weak modules are detected and replaced. The primary risk is not one failed module; it is multiple weak modules, inadequate detection, or a common charging/control problem during a utility outage.

The reliability program therefore focuses on system-level mission reliability. It uses battery monitoring, periodic capacity testing, environmental control, string balancing, charger diagnostics, and spares. Data are segmented by age, temperature, manufacturer batch, and location to identify patterns.

A common software update to all battery-monitoring units is treated as a configuration change with potential common-cause impact. The organization stages deployment and verifies alarms before fleetwide release. Redundancy and condition monitoring are effective only when their dependencies are controlled.

C.8 Municipal water pump motor replacement

A municipal water utility is replacing motors on several high-service pumps. The existing motors are obsolete but have long operating history. A proposed high-efficiency motor offers energy savings, different starting characteristics, and a smaller frame with an adapter base.

The procurement evaluation includes efficiency and first cost but also reliability and maintainability. Engineers examine starting duty, voltage dip, bearing loading, vibration, cooling, spare commonality, alignment, base stiffness, and availability of repair services. A more efficient motor that introduces alignment sensitivity or a proprietary control dependency may not deliver lower lifecycle cost.

The selected design includes standardized interfaces, condition-monitoring provisions, documented alignment tolerances, and a commissioning baseline for vibration and temperature. Reliability requirements are embedded in the procurement rather than left to maintenance after installation.

C.9 Industrial compressor nuisance trips

An industrial compressor has frequent process trips but few mechanical failures. The trip history shows a mix of high-temperature alarms, low suction pressure, and communication-loss events. Maintenance

reports “compressor unreliable,” while the original vendor argues that the machine hardware is functioning correctly.

The system function includes controls and supporting process conditions, so both views are incomplete. Reliability analysis separates mechanical failure, instrument failure, control communication failure, and process-induced protective trips. The largest source of lost availability is a network switch shared with other equipment, followed by an overly sensitive temperature alarm during startup.

Correcting the network architecture and validating alarm setpoints improves availability more than changing compressor hardware. The case demonstrates why reliability boundaries should follow function rather than organizational ownership.

C.10 Elevator fleet recurring door failures

A facility operates a fleet of elevators with frequent door-related service calls. Average availability remains high because technicians respond quickly, but user disruption and maintenance workload are significant. Work orders use broad door-failure codes, limiting analysis.

The reliability team improves failure coding to distinguish sensor alignment, track contamination, operator damage, drive wear, controller faults, and adjustment. Pareto analysis shows that a small subset of high-traffic entrances experiences contamination and repeated sensor misalignment. Cleaning frequency, protective features, sensor mounting, and adjustment standards are changed at those locations.

Fleetwide MTBF improves, but the more important outcome is reduced recurrence at the bad-actor sites. Better data quality enabled a targeted engineering response instead of increasing preventive work across the entire fleet.

C.11 Substation protective relay proof testing

A substation uses protective relays that normally remain passive. Their failure may be hidden until a fault occurs. The maintenance program historically performs time-based proof tests but has gradually extended the interval as digital relays have replaced older electromechanical units.

The engineering review considers relay self-diagnostics, instrument transformers, trip circuits, breaker mechanisms, DC control power, communications, and the consequence of protection failure. Self-diagnostics cover some internal relay failures but do not verify the complete trip path. The proof test therefore remains a functional chain test, with interval and scope based on governing requirements and system consequence.

Test-induced risk is controlled through procedures, configuration checks, independent verification, and post-test restoration. The case illustrates how hidden functions require failure-finding rather than waiting for a real demand.

C.12 Cooling-tower fan gearbox fleet

A plant has twelve cooling-tower fan gearboxes. Four units have experienced bearing failures, all within the same production batch. Other units of different manufacture have accumulated more hours without similar problems. A fleetwide preventive-replacement campaign is proposed.

The reliability review segments the population by manufacturer, batch, lubricant, duty, alignment, and failure mode. Analysis identifies a bearing fit issue in the affected batch. The organization replaces or monitors only that population, increases oil-debris monitoring during the interim, and verifies shaft alignment during the work.

By treating the fleet as several reliability populations rather than one average, the plant avoids unnecessary replacement of healthy gearboxes while addressing the actual common defect.

Worked Reliability Methods - Practical Calculations and Interpretation

D.1 Reliability over a defined mission

Suppose a repairable electronic module is reasonably modeled with a constant failure rate of 2.0 x 10^-5 failures per hour during its useful-life period. For a 1,000-hour mission, the exponential model gives R(t) = exp(-lambda t) = exp(-0.02), or approximately 0.9802. That means the modeled probability of completing the 1,000-hour mission without this failure is about 98.0 percent under the stated assumptions.

The same component over a 10,000-hour mission has R(t) = exp(-0.2), or about 81.9 percent. Nothing changed about the component failure rate; the mission exposure changed. This is why reliability percentages must always include the time or demand basis. Quoting “98 percent reliability” without the mission interval is incomplete.

The example also shows why MTBF should not be treated as a guaranteed life. The reciprocal of the failure rate is 50,000 hours, but the probability of surviving a full 50,000-hour mission under the exponential model would be exp(-1), about 36.8 percent.

D.2 Series-system calculation

Consider a required success path containing a sensor, controller, actuator, and power supply. Their modeled mission reliabilities are 0.995, 0.990, 0.985, and 0.998. If the elements are independent and all four are required, the series reliability is their product, approximately 0.968.

Each component appears highly reliable when viewed individually, yet the complete path has lower reliability because every required element must work. This calculation is useful during architecture review because it can show where redundancy, simplification, better components, or graceful degradation may create the most benefit.

The arithmetic is only as good as the functional model. If the actuator has a manual backup that satisfies the mission requirement, the pure series model may be too pessimistic. If the power supply is actually shared by several redundant trains, the simple component product may be too optimistic. Reliability block diagrams should be reviewed by engineers who understand the physical and operational success criteria.

D.3 Parallel redundancy calculation

Two independent identical pumps each have mission reliability of 0.92, and either pump can satisfy the required function. The idealized parallel reliability is 1 - (1 - 0.92)^2 = 0.9936. The numerical gain is large because the system fails only if both pumps fail during the mission.

Now add a shared suction strainer with mission reliability of 0.97. If that strainer is required by both pumps, the complete success-path reliability cannot exceed 0.97, and the simplified system reliability becomes about 0.9638 when the strainer is modeled in series with the parallel pair. The shared dependency erased much of the apparent redundancy benefit.

This example is intentionally simple, but the principle scales to power, controls, networking, cooling, fuel, environment, software, and human actions. When redundancy is justified by a numerical reliability claim, the shared support paths should be visible in the model.

D.4 Availability calculation and leverage

Assume an asset has MTBF of 2,000 hours and mean restoration time of 20 hours. Using the simple steady-state relationship A = MTBF / (MTBF + MTTR), availability is approximately 99.01 percent. If engineering and maintenance improvements reduce MTTR to five hours without changing failure frequency, availability rises to about 99.75 percent.

Alternatively, if MTTR remains twenty hours but MTBF doubles to 4,000 hours, availability becomes about 99.50 percent. Both paths improve performance, but the faster-restoration project creates the larger availability gain in this example. The calculation helps teams compare where improvement effort may have the best leverage.

Real availability can include planned maintenance, logistics delay, administrative delay, or only corrective downtime depending on the definition. Before comparing reported values, verify which downtime categories are inside the clock.

D.5 Downtime contribution by failure mode

A system experiences three failure modes in one year. Mode A occurs ten times and takes one hour to restore; Mode B occurs twice and takes twelve hours; Mode C occurs once and takes forty hours. If the team prioritizes only by frequency, Mode A appears dominant. If it prioritizes downtime, Mode C contributes forty hours, Mode B twenty-four, and Mode A ten.

This simple calculation illustrates why a Pareto chart of event count can tell a different story from a Pareto chart of downtime, cost, safety exposure, or production loss. Reliability improvement should be prioritized using the consequence measure that aligns with the objective.

A recurring high-frequency nuisance failure may still deserve attention because of labor burden and error exposure, while a rare long-duration failure may dominate availability. Multi-dimensional prioritization is often more useful than one ranked list.

D.6 Preventive replacement threshold thinking

A component shows little age relationship until approximately 12,000 cycles, after which failure frequency rises sharply. The existing preventive replacement interval is 8,000 cycles. Inspection data show that removed components at 8,000 cycles are consistently healthy and that the degradation mechanism is detectable through a measurable dimensional check.

A controlled age-exploration program might extend the interval while increasing targeted inspection near the emerging wear region. The decision should consider the failure consequence, variability in wear, inspection sensitivity, and whether the task restores resistance to the failure mode. The goal is not to “stretch maintenance” but to align intervention with actual failure behavior.

If the failure mode were random and unrelated to cycles, no such threshold logic would exist. A fixed replacement interval could then add cost and maintenance-induced risk without producing meaningful reliability gain.

D.7 P-F interval and monitoring frequency

Suppose a degradation mechanism becomes detectable between four and eight weeks before functional failure, depending on load and environment. A monthly manual inspection may occasionally leave too little warning because the shortest P-F interval is close to the inspection period. A two-week inspection provides better opportunity for detection but increases labor and may still miss rapid degradation.

Continuous or weekly automated monitoring may be justified if consequence is high and the signal is reliable. For lower-consequence equipment, the organization may accept a less aggressive interval combined with a stocked spare and rapid restoration plan. The correct strategy depends on warning-time variability, probability of detection, response lead time, and consequence.

The P-F concept converts condition monitoring from a technology purchase into a timing decision. The sensor or inspection is useful only if the organization can detect, decide, and act before functional failure.

D.8 Spare-part service-level reasoning

A critical drive module has an average failure frequency of once every five years across the fleet, a replacement lead time of nine months, and a two-day installation time if the part is on hand. Without a spare, a failure could create months of outage or require an expensive temporary workaround.

The low annual failure probability can make the spare appear uneconomic if only expected part cost is considered. A reliability decision also considers outage consequence, supplier uncertainty, whether one spare can serve several assets, preservation requirements, obsolescence, and the probability that a stored spare remains usable.

A shared insurance spare may provide substantial risk reduction when interfaces are standardized and transportation is feasible. The engineer should verify logistics before counting the spare as a control: lifting, route clearance, storage environment, firmware, accessories, testing, and installation procedure can all determine whether the spare actually reduces restoration time.

D.9 Repeat-failure rate as an improvement metric

A maintenance organization reports that overall equipment availability improved from 98.8 to 99.2 percent. At the same time, 22 percent of corrective work orders involve a repeat of the same failure mode within ninety days. The headline availability improvement may reflect faster response rather than durable defect elimination.

A repeat-failure metric can reveal this distinction. Definitions should specify the asset boundary, failure mode, and recurrence window. The metric should not punish legitimate repeated wear on low-criticality consumables, but it is useful for identifying incomplete repairs, poor diagnosis, installation error, weak corrective action, or an underlying design problem.

Reliability programs need both restoration and recurrence measures. Restoring service quickly is operationally valuable; preventing the next identical failure is reliability improvement.

D.10 Reliability decision record

A concise reliability decision record should state the required function, operating context, failure mode or reliability issue, evidence, analysis method, assumptions, selected action, residual uncertainty, owner, verification plan, and trigger for reconsideration. This is especially important when the decision departs from a generic preventive interval, accepts run-to-failure, relies on monitoring, or credits redundancy.

The record allows another qualified engineer to reconstruct why the decision was reasonable at the time. It also supports future learning when actual performance differs from the prediction. If a maintenance interval is later found too long, the organization can see which degradation assumptions were wrong. If a spare was never used, the organization can determine whether the risk changed or the original low-frequency/high-consequence logic remains valid.

Documentation should preserve reasoning without becoming paperwork for its own sake. The strongest record is concise, technically specific, and connected to the evidence used to operate and improve the asset.

Design Review Notes - Building Reliability into New Assets

E.1 Functional architecture review

Before detailed design is complete, review the architecture from the perspective of required functions and single points of failure. Identify which functions must remain available after one failure, which can degrade, which can be restored manually, and which require automatic protection. Trace supporting utilities and controls as carefully as primary equipment. A redundant mechanical train that shares one electrical panel or one cooling-water source may not satisfy the intended resilience objective.

Architecture review is the least expensive time to remove hidden dependencies. Later retrofits can require major piping, electrical, software, or structural changes. The review should therefore include operations and maintenance personnel who understand practical access, switching, recovery, and failure history.

E.2 Maintainability walkdown before release

A maintainability walkdown asks whether planned maintenance and likely corrective work can actually be performed in the installed arrangement. Review representative tasks: filter replacement, bearing removal, sensor calibration, valve packing adjustment, software recovery, breaker replacement, lubrication, inspection, and major-component removal. Check clearance, isolation, lifting, drainage, containment, ergonomics, hazardous energy, and the sequence needed to reach the component.

The objective is not to design for every imaginable repair. It is to identify high-frequency and high-consequence tasks whose difficulty would materially increase downtime, safety exposure, or lifecycle cost. Small design changes such as access doors, test points, lifting lugs, removable spools, or connector relocation can have large maintainability value.

E.3 Diagnostic coverage review

For important failure modes, ask how the operator or maintainer will know what failed. A generic alarm that says “fault” may force extensive troubleshooting. Better diagnostic coverage separates sensor fault from process abnormality, local equipment failure from upstream utility loss, and hardware failure from communication or configuration error.

Diagnostics should also support verification after maintenance. The system should make it possible to confirm that the repaired function and its protective features are restored. Logs and trends should preserve enough context to reconstruct intermittent events without requiring the failure to recur under observation.

E.4 Commissioning as reliability evidence

Commissioning and startup provide the first opportunity to verify reliability assumptions under integrated conditions. Functional testing should challenge backup paths, alarms, automatic transfer, degraded modes, protective functions, and recovery procedures rather than only demonstrate nominal operation. Early failures should be investigated because they may reveal installation defects, configuration errors, contamination, poor alignment, or design margin issues.

A commissioning baseline for vibration, temperature, current, pressure, software version, valve position, or other condition measures can improve future monitoring. The baseline establishes what healthy operation looked like before aging and maintenance changed the system.

E.5 Handover to operations and maintenance

Reliability can be lost during turnover if the operating organization does not receive current drawings, spare lists, software backups, calibration data, maintenance procedures, warranty information, configuration records, and training. The handover package should identify critical assets, hidden protective functions, special storage requirements, monitoring thresholds, long-lead spares, and any temporary reliability controls accepted during commissioning.

The operating team should also understand the assumptions behind maintenance intervals and redundancy. When those assumptions are invisible, future staff may disable a backup path, combine previously independent services, or extend a task interval without realizing the reliability consequence. Good handover preserves both equipment and decision knowledge.

E.6 Reliability implementation sequence

A practical reliability program can begin without a large software implementation. First, define the small set of functions and assets whose failure has the greatest consequence. Second, clean the equipment boundary and failure definitions so operating history is interpretable. Third, identify recurring failure modes, hidden protective functions, long restoration times, common dependencies, and maintenance tasks that consume substantial effort without finding degradation. Fourth, select a limited number of improvements with clear technical logic: redesign a recurring mechanism, improve diagnostics, change a

task interval, add a condition indicator, protect a redundant path, stock a justified spare, or improve a restoration plan. Fifth, define how the improvement will be verified.

The program should then expand based on demonstrated value. Reliability work is most credible when operating personnel can see fewer repeat failures, shorter outages, better planned work, safer access, or improved service. Large metric dashboards and asset hierarchies are useful only when they support those outcomes. The engineering discipline is the closed loop: function, evidence, analysis, action, verification, and learning.

For professional engineers, the final test is whether the reliability decision is technically defensible. The analysis should be proportionate to consequence; assumptions should be visible; maintenance should be connected to failure behavior; redundancy should be independent enough to matter; restoration should be feasible; and operating data should be capable of challenging the original assumptions when the asset behaves differently in service.

Principal references and source links

- NASA Reliability-Centered Maintenance Guide for Facilities and Collateral Equipment. Open source
- NASA-STD-8729.1A, June 13, 2017. NASA identifies revision A as active; spaceflight standard used for transferable R&M objectives, not as a general facility mandate. Open source
- NASA Reliability and Maintainability technical discipline, including Physics of Failure resources. Open source
- DOE Operations & Maintenance Best Practices Guide, Release 3.0, August 2010. Figures 6.2.6–6.2.7, printed page 6.6. Open source
- NIST/SEMATECH e-Handbook: exponential reliability model and its constant-hazard assumptions. Open source
- SAE JA1011_202411, Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes. Public title and scope; full licensed text not reproduced. Open source
- ISO 14224:2016, corrected October 2016; published edition, confirmed 2022. Industry-specific reliability and maintenance data framework; public scope used as context. Open source

Actual applications remain subject to the governing project documents, applicable codes and standards, manufacturer instructions, owner requirements and qualified professional judgment. Hypothetical examples identify their own assumptions; source figures retain their provenance and limits. Public standard records were checked November 15, 2026; full licensed ISO and SAE texts were not accessed.

Instructor

Matthew Ward, MBA, CIA

Course scope is limited to general reliability, maintainability, asset-performance and engineering-management concepts. Discipline-specific calculations, equipment diagnosis and design decisions require appropriately qualified professionals.

> **Study guide complete.** Return to the BuilderCE lesson page and select Continue to begin the final examination. A score of at least 70% is required to pass. Retakes are permitted.