

CHAPTER 3

IMPROVING AVAILABILITY OF C4ISR FACILITIES

3-1. Overview of the process

Facility managers are faced with the responsibility of providing the proper utilities (electrical, chilled water, steam, etc.) at the needed levels (power levels, voltage, pressure, etc.) to their customers when needed to support an end mission. The steps for improving the availability of a facility for two situations, new facilities in design and facilities already in use, are shown in table 3-1. Each step for each situation will be discussed in this chapter.

Table 3-1. The process for improving facility availability

New Facilities Being Designed	Facilities Already in Use
1. Determine system availability requirements	1. Determine system availability requirements
2. Derive reliability and maintainability requirements from availability requirement	2. Derive reliability and maintainability requirements from availability requirement
3. Develop "one-lines"	3. Develop "one-lines" of systems
4. Conduct analyses to predict availability, reliability, and maintainability and to determine weaknesses in design and redesign based on failure criteria and cost/benefit analysis	4. Collect data for availability assessment
5. Conduct testing to validate analytical results	5. Assess availability, reliability, maintainability, and logistics performance being achieved for each system (this establishes the baseline performance)
6. Update assessment of availability, reliability, and maintainability based on test results	6. Identify shortfalls (differences between required level of performance and baseline performance)
7. Revise design as necessary based on test results	7. Perform cost-benefit analysis to prioritize improvement efforts
8. Construct facility and continuously assess performance and identify opportunities for improvement	8. Design and develop system changes (using same process used for new facility design)
	9. Assess improvement in availability, reliability, and maintainability based on analyses and test
	10. Implement design changes
	11. Continuously assess performance and identify opportunities for improvement

3-2. New facilities in design

Since reliability and maintainability, and hence availability, are predominantly affected by design, it is essential that these system characteristics be addressed in the design of a new system. It is during design, that these characteristics can be most effectively and positively influenced at the least cost.

a. Determine system availability requirements. Establishing clear, comprehensive, and measurable requirements is the first and most important step in designing and developing systems that meet user needs. The design requirements must be derived from and, if met, allow the user needs to be met. User needs are often stated in non-design terms. For facilities, these might include operational availability, readiness, mean time between maintenance (where maintenance includes all maintenance actions, including those to repair operator-induced failures), and total downtime (including the time to order and ship parts if necessary). Designers must have requirements that they can control. For a facility, these may include inherent availability, mean time between design failures, and mean time to repair (includes only the actual "hands on" time to make a repair). The facility availability requirement should be included in the solicitation package (normally in the specification) for a new facility.

b. Derive reliability and maintainability requirements from availability requirement. Based on the user need (e.g., operational availability), the reliability and maintainability design requirements (e.g., mean time between failure and mean time to repair) must be derived. This derivation of lower-level requirements is usually done by the design organization and continues throughout the development effort until design requirements are available at the lowest level of indenture (subsystem, assembly, subassembly, part) that makes sense.

c. Develop "one-lines". Paragraph 4-5 discusses this method of representing a system.

d. Conduct Analyses. Conduct analyses to predict availability, reliability, and maintainability and to determine weaknesses in design and redesign based on failure criteria and cost/benefit analysis. Some of the pertinent analyses are summarized in table 3-2.

Table 3-2. Analyses helpful in designing for reliability

Analysis	Purpose	Application	When to perform
FEA	<ul style="list-style-type: none"> • Computer simulation technique for predicting material response or behavior of modeled device • Determine material stresses and temperatures • Determine thermal and dynamic loading 	<ul style="list-style-type: none"> • Use for devices that: <ul style="list-style-type: none"> – Are unproven with little prior experience/data – Use advanced/unique packaging/design concepts – Will encounter severe environmental loads – Have critical thermal/mechanical constraints 	In design phase when candidate devices can be selected using selection criteria
TA	<ul style="list-style-type: none"> • Calculate junction temperatures • Calculate thermal gradients • Calculate operating temperatures 	<ul style="list-style-type: none"> • For integrated circuits • For electronics and electrical devices 	<ul style="list-style-type: none"> • During circuit design • Prior to design of cooling systems
Dormancy Analysis	<ul style="list-style-type: none"> • Calculate failure rates of devices while dormant (e.g., storage) 	<ul style="list-style-type: none"> • Use for devices identified to have periods of dormancy 	<ul style="list-style-type: none"> • During design
FTA	<ul style="list-style-type: none"> • Top down approach to identify effects of faults on system safety or reliability • Address multiple failure 	<ul style="list-style-type: none"> • Can be applied when FMECA too expensive • To address effects of multiple failures 	<ul style="list-style-type: none"> • Early in design phase, in lieu of FMECA
FMECA	<ul style="list-style-type: none"> • Bottom up approach to identify single failure points and their effects • To assist in the efficient design of BIT and FIT • To establish and rank critical failures • To identify interface problems 	<ul style="list-style-type: none"> • More beneficial if performed on newly designed equipment • More applicable to equipment performing critical functions (e.g., control systems) 	<ul style="list-style-type: none"> • Early in design phase
SCA	<ul style="list-style-type: none"> • To identify failures not caused by part failures • To reveal unexpected logic flows that can produce undesired results • To expose design oversights that create conditions of undesired operation 	<ul style="list-style-type: none"> • Mission and safety critical functions • Hardware with numerous interfaces • Systems with high testing complexities • Use selectively due to high cost of performing 	<ul style="list-style-type: none"> • Later design stage but prior to CDR
WCCA	<ul style="list-style-type: none"> • To evaluate circuits for tolerance to "drift" • When time dependency is involved • To evaluate the simultaneous existence of all unfavorable tolerances • Single failures 	<ul style="list-style-type: none"> • Assesses combined effect of parts parameters variation and environmental effects on circuit performance • Not often applied • Use selectively 	<ul style="list-style-type: none"> • Later design stage as required

LEGEND: Finite Element Analysis (FEA); Thermal Analysis (TA); Fault Tree Analysis (FTA); Failure Modes, Effects and Criticality Analysis (FMECA); Sneak Circuit Analysis (SCA); Worst Case Circuit Analysis (WCCA)

e. Conduct testing to validate analytical results. No matter how diligent we are in developing the models and analytical tools used to design, we cannot account for all variations and factors. By testing a given design, we will uncover unexpected problems. These problems can include new types of failures, more frequent than expected failures, different effects of failures, and so forth. Problems discovered during test provide opportunities for improving the design and our models and tools.

f. Update assessment of availability, reliability, and maintainability based on test results. Based on the results of our testing, we should update the analytical assessments of reliability made earlier. Adding the results of testing provides higher confidence in our assessment than is possible using analytical results alone.

g. Revise design as necessary based on test results. If our updated assessment indicates we are falling short of our reliability (and availability) requirements, we must revise the design to improve the reliability. Even when our updated assessment indicates that we are or are close to meeting our requirements, we should consider making design changes based on cost-benefit considerations.

h. Construct facility and continuously assess performance and identify opportunities for improvement. Once we are satisfied that the reliability (and availability) requirements are satisfied by our facility design, the facility is constructed. We must ensure that the inherent levels of reliability are sustained over time, and collect information that can be used in the design of the next facility. To that end, we need to collect data and use the data to continuously assess the availability performance of the facility. This operational field data also should be archived for use in designing new facilities.

3-3. Facilities already in use

For facilities in use, the process for improving availability is somewhat different than that discussed for new systems. It is different for two major reasons. First, improvements must be made by modifying an existing design, which is usually more difficult than creating the original design. Second, the improvements must be made with as little disruption to the facility as possible, since it is supporting an ongoing mission. Although design changes are usually the primary focus of improvement efforts, changes in procedures or policy should also be considered. Not only are such changes usually much easier and economical to make, they may actually be more effective in increasing availability.

a. Determine system availability requirements. As was the case for a new system, the requirements must be known. For existing facilities, it may be difficult to find the original user needs or design requirements. Even when the original requirements can be determined, the current requirements may have changed due to mission changes, budget constraints, or other factors.

b. Derive reliability and maintainability requirements from the availability requirement. Whatever the operational requirements are, it is necessary to translate them into reliability and maintainability requirements.

c. Develop "one-lines" of systems. This step can be bypassed if "one-lines" were developed for the facility when it was developed and built and are still current.

d. Collect data for availability assessment. Ideally, a data collection system was implemented when the facility was first put into operation. If that is not the case, one must be developed and implemented at this point. The data to be collected includes failures, failure causes and mechanisms, repair times, and so forth.

e. Assess performance. Assess the availability, reliability, maintainability, and logistics performance being achieved for each system. Performing this step establishes the baseline performance for the facility.

f. Identify shortfalls. Shortfalls are the differences between the required level of performance and baseline performance.

g. Perform cost-benefit analysis to prioritize improvement efforts. Many potential improvements will be identified throughout the life of a facility. Those that are safety-related or are essential for mission success will always be given the highest priority. Others will be prioritized on the basis of the costs to implement compared with the projected benefits. Those that have only a small return for the investment will be given the lowest priority.

h. Design and develop system changes. The process for improving the availability, reliability, and maintainability performance of an existing facility is essentially the same as for designing new facility.

i. Assess improvement. Assess improvement in availability, reliability, and maintainability based on analyses and test. Before implementing any potential improvements, some effort must be made to ensure that the design changes must be validated. All too often, a change that was intended to improve the situation actually makes it worse. Through careful analyses and appropriate testing, one can determine that the proposed change actually results in some level of improvement.

j. Implement design changes. Those design changes that are validated as improving availability must be implemented in a way that minimizes the downtime of the facility. Perhaps they can be made during scheduled maintenance periods. Or perhaps there are times of the day, month, or year when downtime is less critical to the mission than at other times. Careful planning can minimize the impact on the mission. Also, the procedures, tools, training, and materials needed for the design change must be in place and validated prior to starting the facility modification.

k. Monitor performance. Continuously assess performance and identify opportunities for improvement. Continuous improvement should be the goal of every facility manager. As the facility ages, the cost-benefits of what were low-priority improvements may change, new problems may be introduced, and new mission requirements may arise. By collecting data and maintaining a baseline of the facility availability performance, the facility manager will be in a position to make future improvements as they become necessary or economical.

3-4. Improving availability through addition of redundancy

Redundancy is a technique for increasing system reliability and availability by making the system immune to the failure of a single component. It is a form of fault tolerance – the system can tolerate one or more component failures and still perform its function(s).

a. *Types of Redundancy.* There are essentially two kinds of redundancy techniques employed in fault tolerant designs, space redundancy and time redundancy. Space redundancy provides separate physical copies of a resource, function, or data item. Time redundancy, used primarily in digital systems, involves the process of storing information to handle transients, or encoding information that is shifted in time to check for unwanted changes. Space, or hardware, redundancy is the approach most commonly associated with fault tolerant design. Figure 3-1 provides a simplified tree-structure showing the various types of hardware redundancy that have been used or considered in the past.

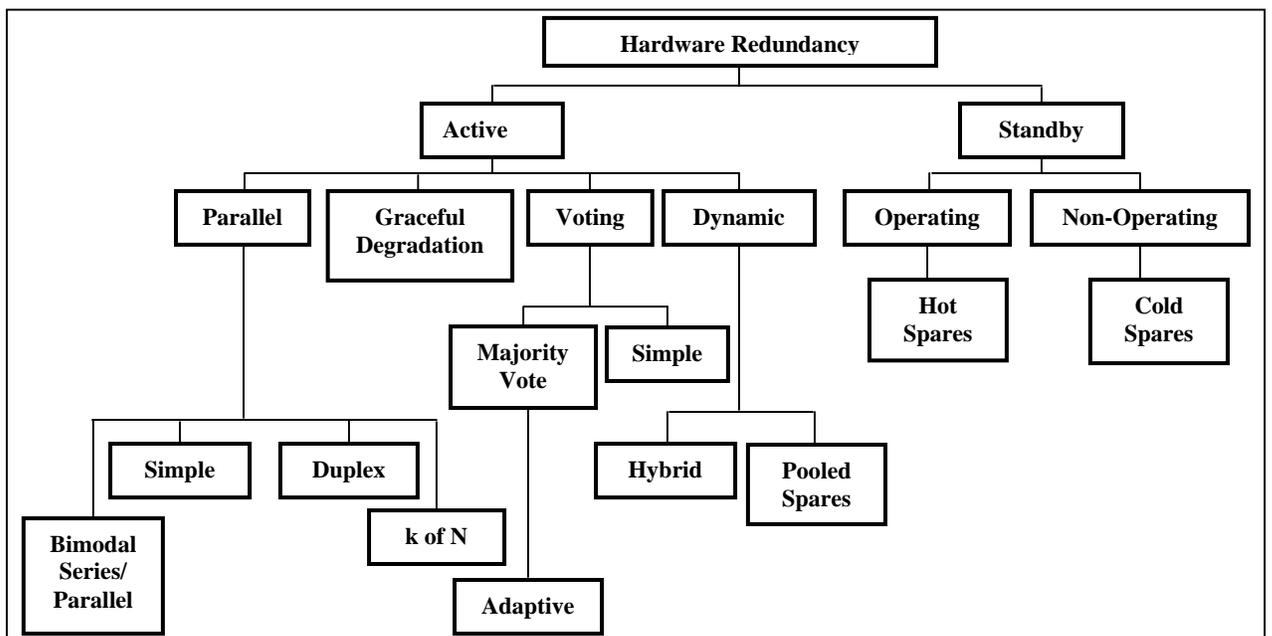


Figure 3-1. Types of redundancy.

b. *Impact on Testability.* Many of today's more sophisticated systems not only require an ability to detect faults but also to diagnose or isolate them. It may even be desirable for a system to have the ability to reconfigure itself to avoid system failure. Automated fault detection and isolation has therefore become an essential means of obtaining highly fault-tolerant systems. Because of this, the design of the diagnostic system, including any built-in-test (BIT) features and the overall testability of the design are important tradeoffs that need to be made as part of the fault tolerant design process. Table 3-3 presents a sample list of hardware fault tolerant design approaches, and their impact on diagnostic approaches and BIT.

Table 3-3. Diagnostic implications of fault tolerant design approaches

Fault Tolerant Design Technique	Description	Diagnostic Design Implications	BIT Implications
Active Redundancy, simple parallel	All parallel units are on whenever the system is operating. k of the N units are needed, where $0 < k < N$. External components are not required to perform the function of detection, decision and switching when an element or path in the structure fails. Since the redundant units are always operating, they automatically pick up the load for a failed unit. An example is a multi-engined aircraft. The aircraft can continue to fly with one or more engines out of operation.	Hardware/Software is more readily available to perform multiple functions.	N/A
Active Redundancy with voting logic	Same as Active Redundancy but where a majority of units must agree (for example, when multiple computers are used)	Performance/status-monitoring function assures the operator that the equipment is working properly; failure is more easily isolated to the locked-out branch by the voting logic.	N/A
Stand-by redundancy (Non-operating)	The redundant units are not operating and must be started if a failure is detected in the active unit (e.g., a spare radio is turned on when the primary radio fails).	Test capability and diagnostic functions must be designed into each redundant or substitute functional path (on-line AND off-line) to determine their status.	Passive, periodic, or manually initiated BIT.
Stand-by redundancy (Operating)	The redundant units are operating but not active in system operation; must be switched “in” if a failure is detected in the active unit (e.g., a redundant radar transmitter feeding a dummy load is switched into the antenna when the main transmitter fails).	N/A	Limited to passive BIT (i.e., continuous monitoring) supplemented with periodic BIT.

(1) No matter which technique is chosen to implement fault tolerance in a design, the ability to achieve fault tolerance is becoming increasingly dependent on the ability to detect, and isolate malfunctions as they occur or are anticipated to occur. Alternate maintainability diagnostic concepts must be carefully reviewed for effectiveness before committing to a final design approach. In particular, BIT design has become very important to achieving a fault tolerant system. When using BIT in fault tolerant system design, the BIT system must: do the following.

- (a) Maintain real-time status of the system’s assets (on-line and off-line, or standby, equipment).
- (b) Provide the operator with the status of available system assets.
- (c) Maintain a record of hardware faults for post-mission evaluation and corrective maintenance.

(2) The essence of fault tolerance is that the system is able to perform its mission despite experiencing some failures. In systems where redundancy is used, this fault tolerance is achieved by one or more redundant units taking over the function previously being performed by another unit. When standby redundancy is used, the failed unit must be detected and the standby unit “brought on line.” In still other cases, principally involving electronics, failures can be “repaired” by rerouting signals or functions to other units. These “repairs” can be done upon a failure or in anticipation of a failure. In such cases, the BIT should, in addition to the actions identified in paragraph 3-4b(1), maintain a record of any reconfiguration events that were required for system recovery during the mission.

(3) For fault tolerant systems, it is important that the design’s inherent testability provisions include the ability to detect, identify, recover, and if possible reconfigure, and report equipment malfunctions to operational personnel. The reliability block diagrams for fault tolerant systems are complex, with non-serial connections. Fault tolerant systems often have a multitude of backups with non-zero switch-over time and imperfect fault detection, isolation, and recovery. Therefore, it is imperative that effective testability provisions be incorporated in the system design concept. If they are not, the fielded design will exhibit long troubleshooting times, high false alarm rates, and low levels of system readiness.

c. Reliability's role in the fault tolerant design process. The role of the reliability engineer in regards to fault tolerant design requirements is to ensure that system reliability requirements are achievable for each of the fault tolerant design approaches being considered. Furthermore, to properly design a fault tolerant system, including a diagnostic scheme, the designer needs to understand the modes in which the system can fail, and the effects of those failure modes. This requires that a failure mode and effects analysis (FMEA) be performed, as a minimum. The FMEA will identify which faults can lead to system failure and therefore must be detected, isolated and removed to maintain system integrity. In general, the reliability design manager must ask a series of questions, as listed in table 3-4.

d. Fault tolerance and tradeoffs. The designer needs to consider each of the questions in table 3-4 and others as part of the overall fault tolerant design process. Other reliability tradeoffs to be considered involve analysis of the redundancy approaches being considered for the fault tolerant design. In addition to reliability concerns, fault tolerance also requires analysis of the impacts on maintainability and testability. As an example, consider figure 3-2. This figure illustrates a design vs. corrective maintenance tradeoff analysis performed early in the product development phase. In particular, the figure shows the tradeoff of restoration frequency versus the number of sensors being used to meet requirements. This program requires a time period for allocating a scheduled maintenance activity and a probability of less than one in 10 billion per flight hour that a total loss of the skewed sensor function would occur. The tradeoff is made between the number of sensors and the cost of unscheduled maintenance activity associated with each approach. Other tradeoffs, such as cost, power, weight, etc. are also necessary. In general, as in any design analysis support function, an analysis of the impacts on reliability, maintainability (including testability) and availability of a chosen fault tolerant design approach must be performed.

Table 3-4. Questions for the reliability design engineer related to fault tolerance

1. How do the system fault tolerance requirements impact the overall reliability, maintainability, and availability requirements?
2. Where should fault tolerant design methods be applied?
 - Which functions involve the most risk to mission success?
 - What is the effect of the operating environment
 - What maintenance strategy/policy needs to be considered?
3. What is the effect on maintainability and testability?
4. What are the constraints that affect fault tolerance?
 - Cost
 - Size & weight
 - Power
 - Interface complexity
 - Diagnostic uncertainties

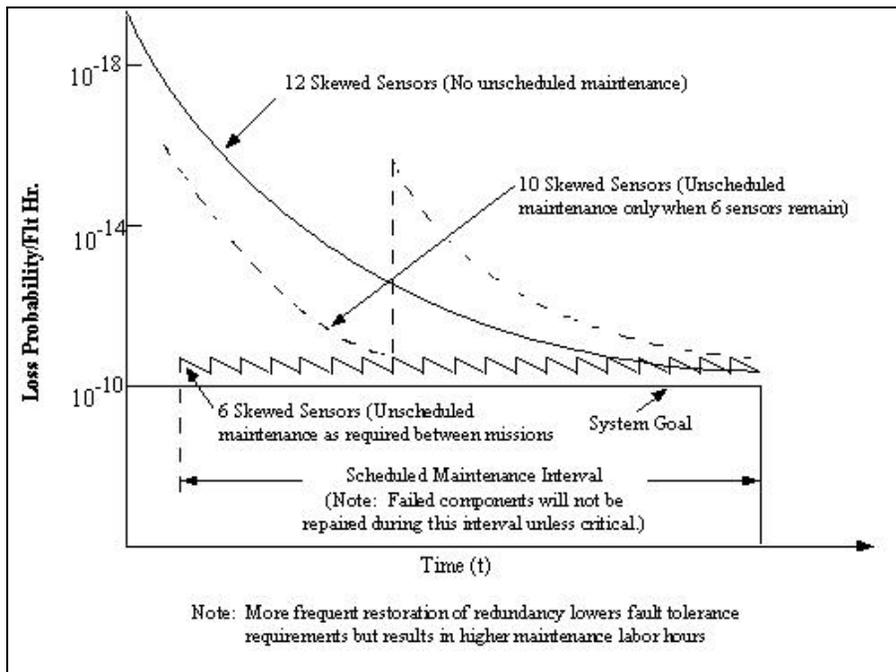


Figure 3-2. Effect of maintenance concept on level of fault tolerance.

e. *General rules in applying redundancy.* In applying redundancy to a C4ISR facility, the following general rules should be followed:

(1) Determine the weak links in the system to know where to add redundancy. These weak links may be portions of the system prone to single point failures or, where redundancy is already used, the reliability is still too low to meet availability requirements.

(a) As an example of applying rule (1), consider the simple system shown in figure 3-3. This system has five subsystems (lettered) with seven major components (numbered). The MTBF and MTTR for each component are shown. Using these figures, the overall system availability can be calculated using Monte Carlo simulation. The results of a Monte Carlo simulation of the system using RAPTOR yielded the results shown in table 3-5. The areas of weakness from a availability perspective can be determined from simply looking at the relative contribution to system unreliability as summarized in table 3-6. Note that subsystem C is the weakest link, even though it is not subject to a single point failure. Subsystem D is the next weakest link; it is subject to a single point failure. It may have been obvious that D, representing

a potential single point failure, is a weak link. It may not have been as obvious that C, even though it already incorporates redundancy, is a weak point. Looking at the relative availability of component 3, we see that it is much less reliable than the other components. Even dual redundancy is insufficient to compensate for the low MTBF. As this example shows, although it may be tempting to always add redundancy to those portions of a system subject to single point failures, it is sometimes more effective to add it elsewhere.

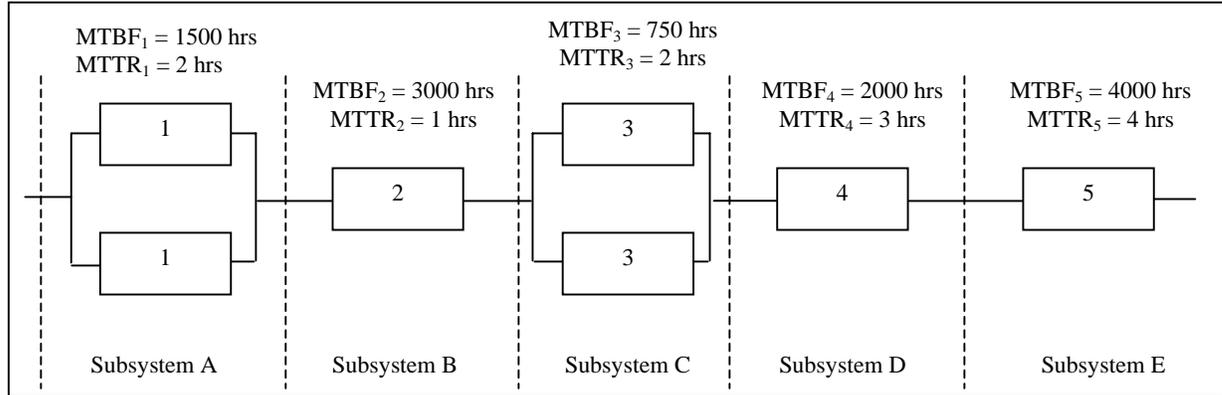


Figure 3-3. Analyzing the contributions to system reliability helps determine where redundancy is needed.

Table 3-5. Calculated availability of system in figure 3-3 using RAPTOR.

MTBM	Mean System Failures	MTTR	Availability (%)
258.77	1.0658	2.5695	99.7236

- Notes:
1. For ease of calculation, the times to failure and the times to repair were assumed to be distributed exponentially.
 2. 10,000 simulation trials were run using an operating time of 1,000 hours.

Table 3-6. Relative unreliability of subsystems (repairs ignored)

Subsystem	Reliability in 1000 hours	Expected Failures per 1000 Hours	% Contribution to System Unreliability	Contribution to System Unreliability Ranking
A	0.7632	0.2368	14.12	4
B	0.7165	0.2835	16.90	3
C	0.4577	0.5423	32.33	1
D	0.6065	0.3935	23.46	2
E	0.7788	0.2212	13.19	5
SYSTEM	0.1182	1.6773	-	-

(2) Add redundancy in a way that avoids undesirable interactions. Rule 2 implies that some components cannot be used in some forms of redundancy, depending on the failure modes, application, and other factors. The type of redundancy shown in figure 3-3 is active redundancy, in which all components are on all of the time that the system is operating. In some cases, such a redundant configuration would result in undesired interactions or interference among the redundant units. As will be seen later in this chapter, certain forms of redundancy are preferable to others in a given application.

(3) Adding redundancy increases support requirements and costs. Only use redundancy when availability is insufficient and no other technique will improve it. Rule 3 refers to the added costs incurred with redundancy. The most obvious increase is due to the fact that more components must be purchased and installed. An additional cost comes from an increased failure rate. The increase in complexity results in an increase in unscheduled maintenance. If nothing is done to

improve the reliability of the individual components in a system, but additional components are added to provide redundancy, the total failure rate of the components will increase. System reliability will improve but more component failures will occur. These failures will increase support requirements and costs. Redundancy also increases weight, space requirements, complexity, and time to design. Thus, safety and mission reliability is gained at the expense of adding an item(s) in the unscheduled maintenance chain. Only use redundancy when availability is insufficient and no other technique will improve it.

(a) The decision to use redundant design techniques must be based on analysis of the tradeoffs involved. Redundancy may prove to be the only available method, when other techniques of improving reliability, e.g., derating, simplification, better components, have been exhausted, or when methods of item improvement are shown to be more costly than duplications.

(b) When preventive maintenance is planned, the use of redundant equipment can allow for repair with no system downtime. Occasionally, situations exist in which equipments cannot be maintained, e.g., satellites; then redundant elements may be the best way to significantly prolong operating time.

(4) Ensure that any one redundant unit can be maintained without shutting down the other redundant units. Rule 4 requires that we ensure that any one redundant unit can be maintained without shutting down the other redundant units. Assume that two generators, for example, are sharing a load. If one fails and we must shut the other generator down to either gain access to or repair the failed generator, then we in effect have no redundancy. An implicit assumption in using redundancy is that availability increases because we can repair a failed component while the remaining redundant components continue to operate. If this assumption is violated, redundancy will not increase availability.

f. Design considerations. The FMEA is a primary reliability analysis, critical to the fault tolerant design process. The reliability engineer will also use additional techniques for analyzing a fault tolerant design to verify that it meets reliability requirements. However, many of the evaluation tools used in the past are no longer adequate to deal with more sophisticated fault tolerant designs that include more complex fault handling capabilities. Because fault handling methods include the use of fault detection and fault recovery approaches, any evaluation tool must include the ability to properly account for the effects of imperfect fault coverage (or fault detection) and fault recovery.

(1) Monte Carlo simulation and Markov techniques continue to be used as the primary means of analyzing highly sophisticated fault tolerant designs. These approaches have been modified to incorporate situations where the sequence of failure is important, where the failure is transient or intermittent, or where the response to failure (i.e., detection, isolation, recovery, and reconfiguration) is imperfect. In these situations, Markov methods continue to lead the way in evaluation methods. In general, the Markov approach, which is used to define the specific states that a system can occupy, has been used to incorporate fault handling and recovery. A major limitation to the Markov approach is that the number of system states that must be defined to comprehensively describe a large system and model the behavior of complex fault management schemes can become very large (approaching 10^5 for highly complex systems). A common solution to this problem is to partition the system into smaller systems, evaluate each partition separately, and then combine the results at the system level. However, such an approach is only exact when each partitioned subsystem's fault tolerant behavior is mutually independent of each other. If subsystem dependencies do exist, then an assumption of independence will result in only an approximate solution.

(2) Other approaches that are now becoming more common involve decomposing the system into separate fault-occurrence and fault handling submodels. However, the inputs for this type of approach require knowledge of the distribution and parameter values of: detection, isolation, recovery, rates, etc. The following is a list of assumptions, limitations and sources of error found in existing reliability models:

(a) Solving a fault-handling model in isolation and then reflecting its results in an aggregate model is, itself, an approximation technique. The assumptions necessary to determine a solution typically result in a lower bound (conservative) approximation of the system reliability.

(b) Separate fault-handling models have been assumed to be independent of system state. This requires that the same fault-handling model and choice of parameters be used irrespective of the system's level of degradation. This ignores the fact that for many systems the recovery process is faster if the number of active units is smaller or that the recovery process may be different, depending on the sequence of events in different subsystems.

(c) The common technique of partitioning the system into independent functional subgroups for computational ease is a potential source of error. The magnitude and direction of the error is a function of how truly independent/dependent the subgroups are of each other. If subgroups are assumed independent when in fact they are not, the effect is an overstatement of system reliability/availability. If subgroups are assumed completely dependent when some degree of independence exists, the effect is an understatement of the system's reliability/availability.

(d) Some models assume a constant instantaneous fault-protection coverage factor in lieu of a separate fault handling model. These fail to recognize that during time spent in the intermediate fault-handling states to detect, isolate, and recover/reconfigure, a second item failure could result in system failure. Further, as with fault handling models, these times are generally not constant, but depend on the current state of the system.

(e) Most models require the assumption that the system is perfect at the mission start. Therefore, they cannot evaluate the effects of latent defects (e.g., handling, manufacturing, transportation, and prior mission), nor assist in determining the testability payoff or requirements for detection and removing them before the start of the mission. Models with this limitation cannot be used to evaluate alternate maintenance concepts that include degradation between missions as an acceptable strategy.

(f) Some models require that spares be treated exactly like active units, irrespective of their actual utilization in the system mechanization. This requires that spares are assumed to be "hot" and have the same failure rates and failure modes as the active units. This assumption will cause the model to understate the system reliability in those situations where spares are "cold" or in "stand-by" and/or where their failure rates may be less than those of the active units.

(g) As indicated previously, some models require the assumption that item failure rates are constant throughout time. This will result in an overstatement of system reliability if the items have failure rates that increase with mission time. Some models remove this restriction and permit time-varying failure rates. However, the solution the algorithms employ requires the use of global time (as opposed to local time of entry into a state), thus precluding the use of the model for repairable systems and availability analysis.

3-5. Improving availability through reliability-centered maintenance (RCM)

All C4ISR facilities that are currently in operation require maintenance to continue to properly perform their functions and support their assigned missions. An effective and efficient maintenance program saves resources and maximizes availability. Reliability-Centered Maintenance (RCM) is an approach for developing an effective and efficient maintenance program based on the reliability characteristics of the constituent parts and subsystems, economics, and safety.

a. RCM introduction. Prior to the development of the RCM methodology, it was widely believed that everything had a "right" time for some form of preventive maintenance (PM), usually replacement or overhaul. Despite this commonly accepted view, the results indicated that in far too many instances, PM seemed to have no beneficial effects, and, in many cases, actually made things worse by providing more opportunity for maintenance-induced failures.

b. RCM overview. The RCM approach provides a logical way of determining if PM makes sense for a given item and, if so, selecting the appropriate type of PM. The approach is based on:

(1) RCM seeks to preserve system or equipment function, not just operability for operability's sake.

(2) RCM is more concerned on maintaining end system function than individual component function.

(3) Use reliability as the basis for decisions. The failure characteristics of the item in question must be understood to determine the efficacy of preventive maintenance.

(4) Consider safety first and then economics. Safety must always be preserved. When safety is not an issue, preventive maintenance must be justified on economic grounds.

(5) Acknowledge design limitations. Maintenance cannot improve the inherent reliability – it is dictated by design

(6) Treat RCM as a continuing process. The difference between the perceived and actual design life and failure characteristics is addressed through age (or life) exploration.

c. Preventive maintenance. RCM has changed the approach to preventive maintenance. The RCM concept has completely changed the way in which PM is viewed. It is now widely accepted that not all items benefit from PM, and it is often less expensive (in all senses of that word) to allow an item to "run to failure" rather than to do PM.

d. RCM definitions. The following definitions are commonly used in connection with RCM.

(1) RCM is a logical, structured framework for determining the optimum mix of applicable and effective maintenance activities needed to sustain the operational reliability of systems and equipment while ensuring their safe and economical operation and support.

(2) Maintenance is defined as those activities and actions that directly retain the proper operation of an item or restore that operation when it is interrupted by failure or some other anomaly. (Within the context of RCM, proper operation of an item means that the item can perform its intended function.

(3) Corrective maintenance is maintenance required to restore a failed item to proper operation. Restoration is accomplished by removing the failed item and replacing it with a new item, or by fixing the item by removing and replacing internal components or by some other repair action.

(4) Scheduled and Condition-based preventive maintenance conducted to ensure safety, reduce the likelihood of operational failures, and obtain as much useful life as possible from an item

e. Condition monitoring and analysis. Some impending failures can be detected using some form of condition monitoring and analysis, a type of preventive maintenance. Condition monitoring is defined as periodically or continuously checking physical characteristics or operating parameters of an item. Based on analyzing the results of condition monitoring, a decision is made to either take no action or to replace or repair the item. Condition monitoring can be performed through inspection, or by monitoring performance or other parameters.

f. The RCM concept. RCM has two primary objectives: to ensure safety through preventive maintenance actions, and, when safety is not a concern, preserve functionality in the most economical manner. Preventive Maintenance (PM) is applicable only if it is both effective and economically viable. When safety is not a consideration and PM is either not effective or less economical than running to failure, only CM is required.

(1) PM can be effective only when there is a quantitative indication of an impending functional failure or indication of a hidden failure. That is, if reduced resistance to failure can be detected (potential failure) and there is a consistent or predictable interval between potential failure and functional failure, then PM is applicable.

(2) . The costs incurred with any PM being considered for an item must be less than for running the item to failure (economic viability). The failure may have operational or non-operational consequences. The two categories of cost included in such a comparison for these two failure consequences are (1) operational - the indirect economic loss as a result of failure and the direct cost of repair, and (2) non-operational - the direct cost of repair.

g. A product can fail in two basic ways. First, it can fail to perform one or more of the functions for which it was designed. Such a failure is called a functional failure. Second, a product can fail in such a way that no function is impaired. The failure could be something as simple as a scratch or other damage of the finish of the product. Or it could be that one of two redundant items, only one of which is required for a given function, has failed.

h. The three categories of failure consequences generally used in RCM analysis are Safety, Operational, and Economic. If a functional failure directly has an adverse affect on operating safety, the failure effect is categorized as Safety. When the failure does not adversely affect safety but prevents the end system from completing a mission, the failure is categorized as an Operational failure. When a functional failure does not adversely affect safety and does not adversely affect operational requirements, then the failure is said to have an Economic effect. The only penalty of such a failure is the cost to repair the failure.

3-6. Application of RCM to C4ISR facilities

For equipment used in facilities, condition monitoring, including inspections, overhauls, lubrication and servicing, and failure-finding tasks are all routinely part of an RCM-based

preventive maintenance program. C4ISR facilities potentially require all these tasks. More detailed information on applying RCM to C4ISR facilities will appear in TM 5-698-2 when written.