

2013

Aviation Maintenance Management

Harry A. Kinnison

Tariq Siddiqui
tariq836@erau.edu

Follow this and additional works at: <https://commons.erau.edu/publication>



Part of the [Aviation Safety and Security Commons](#), and the [Maintenance Technology Commons](#)

Scholarly Commons Citation

Kinnison, H. A., & Siddiqui, T. (2013). Aviation Maintenance Management. *Aviation Maintenance Management*, (Second Edition). Retrieved from <https://commons.erau.edu/publication/1538>

This Book is brought to you for free and open access by Scholarly Commons. It has been accepted for inclusion in Publications by an authorized administrator of Scholarly Commons. For more information, please contact commons@erau.edu.

AVIATION MAINTENANCE MANAGEMENT

2ND EDITION



HARRY A. KINNISON

TARIQ SIDDIQUI

Fundamentals of Maintenance

"... maintenance is a science since its execution relies, sooner or later, on most or all of the sciences. It is an art because seemingly identical problems regularly demand and receive varying approaches and actions and because some managers, foremen, and mechanics display greater aptitude for it than others show or even attain. It is above all a philosophy because it is a discipline that can be applied intensively, modestly, or not at all, depending upon a wide range of variables that frequently transcend more immediate and obvious solutions."

LINDLEY R. HIGGINS
Maintenance Engineering Handbook;
McGraw-Hill, NY, 1990.

These opening chapters contain basic information related to the aviation maintenance field and should be considered background for the maintenance management effort. Chapter 1 begins with a discussion of the fundamental reasons why we have to do maintenance in the first place. After all, our skills and techniques have improved immensely over the 100-year history of flight, but we haven't quite reached total perfection. And, considering the number of components on a modern aircraft, we realize early on that maintenance is a complex, ongoing process. For that reason, we need to approach it systematically.

We need a well-thought-out program to address the diverse activities we will encounter in this endeavor, so in Chap. 2 we will study the industry procedures for developing an initial maintenance program. We will discuss the various maintenance check packages (the 48-hour and transit check, the monthly "A" check, the yearly "C" check, etc.)

2 Fundamentals of Maintenance

used to implement the maintenance tasks. We then address the ongoing process of adjusting that program during the lifetime of the equipment. In Chap. 3, we establish the goals and objectives for an airline maintenance program that will serve the real-life operation.

Chapter 4 discusses the extensive certification requirements levied on the aviation industry from the original design of the vehicle to the establishment of commercial operators and the people who run them. The documentation for the aircraft, its operation, and its maintenance, is discussed in Chap. 5 and includes the documents produced by the equipment manufacturers, by the regulatory authorities, and by the airline itself.

Chapter 6 will identify those activities required by the FAA to accomplish maintenance as well as those additional requirements deemed necessary by operators to coordinate and implement an effective maintenance and engineering program. Chapter 7 defines a maintenance and engineering (M&E) organization for a typical midsized airline. Variations for larger and smaller airlines will also be discussed. Part I, then, can serve as background to the remainder of the book and can, if desired, be used as the basis for a first or introductory course on the subject of aviation maintenance management.

Why We Have to Do Maintenance

Introduction

Why do we have to do maintenance? It is simple: “The maintenance of an aircraft provides assurance of flight safety, reliability, and airworthiness.” The aircraft maintenance department is responsible for accomplishing all maintenance tasks as per the aircraft manufacturer and the company’s requirements. The goal is a safe, reliable, and airworthy aircraft.

The aircraft maintenance department provides maintenance and preventive maintenance to ensure reliability, which translates into aircraft availability. These functions do not preclude a random failure or degradation of any part or system, but routine maintenance and checks will keep these from happening and keep aircraft in good flying condition.

Thermodynamics Revisited

Nearly all engineering students have to take a course in thermodynamics in their undergraduate years. To some students, aerodynamicists and power plant engineers for example, thermodynamics is a major requirement for graduation. Others, such as electrical engineers for instance, take the course as a necessary requirement for graduation. Of course, thermodynamics and numerous other courses are “required” for all engineers because these courses apply to the various theories of science and engineering that must be understood to effectively apply the “college learning” to the real world. After all, that is what engineering is all about—bridging the gap between theory and reality.

There is one concept in thermodynamics that often puzzles students. That concept is labeled *entropy*. The academic experts in the thermodynamics field got together one day (as one thermo professor explained) to create a classical thermodynamic equation describing all the energy of a system—any system. When they finished, they had an equation of more than several terms; and all

4 Fundamentals of Maintenance

but one of these terms were easily explainable. They identified the terms for heat energy, potential energy, kinetic energy, etc., but one term remained. They were puzzled about the meaning of this term. They knew they had done the work correctly; the term had to represent energy. So, after considerable pondering by these experts, the mysterious term was dubbed “unavailable energy”—energy that is unavailable for use. This explanation satisfied the basic law of thermodynamics that energy can neither be created nor destroyed; it can only be transformed. And it helped to validate their equation.

Let us shed a little more light on this. Energy is applied to create a system by manipulating, processing, and organizing various elements of the universe. More energy is applied to make the system do its prescribed job. And whenever the system is operated, the sum total of its output energy is less than the total energy input. While some of this can be attributed to heat loss through friction and other similar, traceable actions, there is still an imbalance of energy. Defining entropy as the “unavailable energy” of a system rectifies that imbalance.

The late Dr. Isaac Asimov, biophysicist and prolific writer of science fact and science fiction,¹ had the unique ability to explain the most difficult science to the layperson in simple, understandable terms. Dr. Asimov says that if you want to understand the concept of entropy in practical terms, think of it as the difference between the theoretically perfect system you have on the drawing board and the actual, physical system you have in hand. In other words, we can design perfect systems on paper, but we cannot build perfect systems in the real world. The difference between that which we design and that which we can build constitutes the natural entropy of the system.

A Saw Blade Has Width

This concept of entropy, or unavailable energy, can be illustrated by a simple example. Mathematically, it is possible to take a half of a number repeatedly forever. That is, half of one is $1/2$; half of that is $1/4$, half of that is $1/8$, and so on to infinity. Although the resulting number is smaller and smaller each time you divide, you can continue the process as long as you can stand to do so, and you will never reach the end.

Now, take a piece of wood about 2 feet long (a 2×4 will do) and a crosscut saw. Cut the board in half (on the short dimension). Then take one of the pieces and cut that in half. You can continue this until you reach a point where you can no longer hold the board to saw it. But, even if you could find some way to hold it while you sawed, you would soon reach a point where the piece you have left to cut is thinner than the saw blade itself. When (if) you saw it one more time, there will be nothing left at all—nothing but the pile of sawdust on the floor. The number of cuts made will be far less than the infinite number of times that you divided the number by two in theory.

¹Dr. Asimov wrote over 400 books during his lifetime.

The fact that the saw blade has width and that the act of sawing creates a kerf in the wood wider than the saw blade itself, constitutes the entropy of this system. And no matter how thin you make the saw blade, the fact that it has width will limit the number of cuts that can be made. Even a laser beam has width. This is a rather simple example, but you can see that the real world is not the same as the theoretical one that scientists and some engineers live in. Nothing is perfect.

The Role of the Engineer

The design of systems or components is not only limited by the imperfections of the physical world (i.e., the “natural entropy” of the system), it is also limited by a number of other constraints which we could refer to as “man-made entropy.” A design engineer may be limited from making the perfect design by the technology or the state of the art within any facet of the design effort. He or she may be limited by ability or technique; or, more often than not, the designer may be limited by economics; that is, there just is not enough money to build that nearly perfect system that is on the drawing board or in the designer’s mind. Although the designer is limited by many factors, in the tradition of good engineering practice, the designer is obliged to build the best system possible within the constraints given.

Another common situation in design occurs when the designer has produced what he or she believes is the optimum system when the boss, who is responsible for budget asks, “How much will it cost to build this?” The designer has meticulously calculated that these widgets can be mass produced for \$1200 each. “Great,” says the boss. “Now redesign it so we can build it for under a thousand dollars.” That means redesign, usually with reduced tolerances, cheaper materials, and, unfortunately, more entropy. More entropy sometimes translates into more maintenance required. The design engineer’s primary concern, then, is to minimize (not eliminate) the entropy of the system he or she is designing while staying within the required constraints.

The Role of the Mechanic

The mechanic [aircraft maintenance technician (AMT), repairer, or maintainer], on the other hand, has a different problem. Let us, once again, refer to the field of thermodynamics. One important point to understand is that entropy not only exists in every system, but that the entropy of a system is always increasing. That means that the designed-in level of perfection (imperfection?) will not be permanent. Some components or systems will deteriorate from use, and some will deteriorate from lack of use (time or environment related). Misuse by an operator or user may also cause some premature deterioration or degradation of the system or even outright damage. This deterioration or degradation of the system represents an increase in the total entropy of the system. Therefore, while the engineer’s job is to minimize the

6 Fundamentals of Maintenance

entropy of a system during design, the mechanic's job is to combat the natural, continual increase in the entropy of the system during its operational life-time.

To summarize, it is the engineer's responsibility to design the system with as high degree of perfection (low entropy) as possible within reasonable limits. The mechanic's responsibility is to remove and replace parts, troubleshoot systems, isolate faults in systems by following the fault isolation manual (FIM, discussed in Chap. 4), and restore systems for their intended use.

Two Types of Maintenance

Figure 1-1 is a graph showing the level of perfection of a typical system. One hundred percent perfection is at the very top of the y-axis. The x-axis depicts time. There are no numbers on the scales on either axis since actual values have no meaning in this theoretical discussion. The left end of the curve shows the level of perfection attained by the designers of our real world system. Note that the curve begins to turn downward with time. This is a representation of the natural increase in entropy of the system—the natural

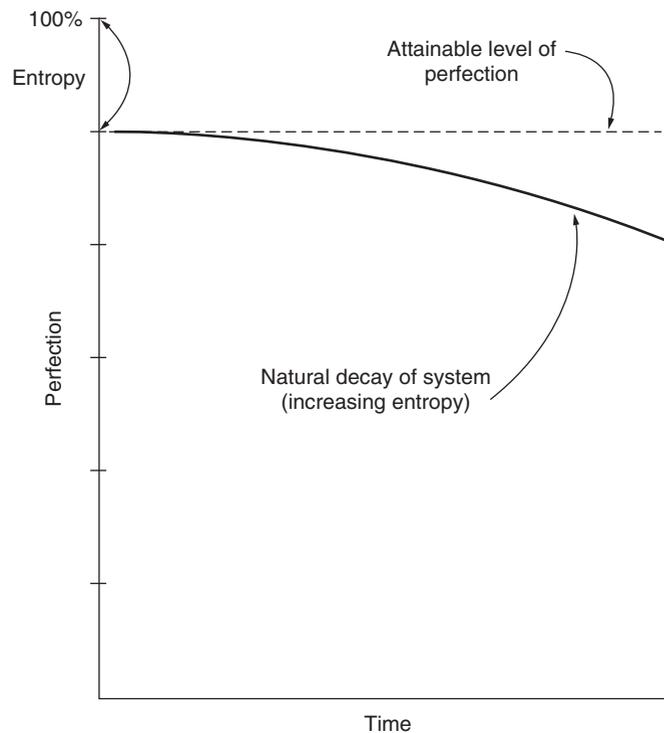


Figure 1-1 The difference between theory and practice.

deterioration of the system—over time. When the system deteriorates to some lower (arbitrarily set) level of perfection, we perform some corrective action: adjusting, tweaking, servicing, or some other form of maintenance to restore the system to its designed-in level of perfection. That is, we reduce the entropy to its original level. This is called *preventive maintenance* and is usually performed at regular intervals. This is done to prevent deterioration of the system to an unusable level and to keep it in operational condition. It is sometimes referred to as *scheduled maintenance*. This schedule could be daily, every flight, every 200 flight hours, or every 100 cycles (a cycle is a takeoff and a landing).

Figure 1-2 shows the system restored to its normal level (curves a and b). There are times, of course, when the system deteriorates rather rapidly in service to a low level of perfection (curve c). At other times the system breaks down completely (curve d). In these cases, the maintenance actions necessary to restore the system are more definitive, often requiring extensive testing, troubleshooting, adjusting, and, very often, the replacement, restoration, or complete overhaul of parts or subsystems. Since these breakdowns occur at various, unpredictable intervals, the maintenance actions employed to correct the problem are referred to as *unscheduled maintenance*.

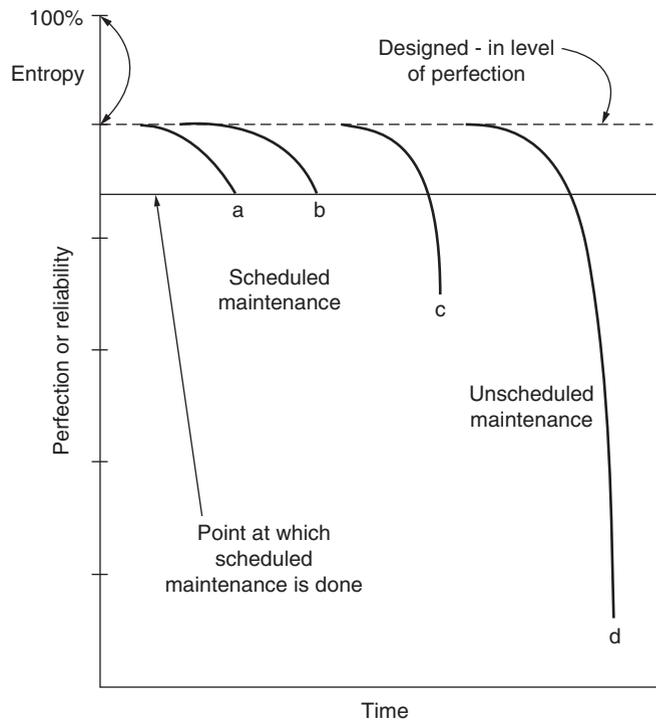


Figure 1-2 Restoration of system perfection.

8 Fundamentals of Maintenance

Reliability

The level of perfection we have been talking about can also be referred to as the reliability of the system. The designed-in level of perfection is known as the inherent reliability of that system. This is as good as the system gets during real world operation. No amount of maintenance can increase system reliability any higher than this inherent level. However, it is desirable for the operator to maintain this level of reliability (or this level of perfection) at all times. We will discuss reliability and maintenance in more detail in Chap. 19. But there is one more important point to cover—redesign of the equipment.

Redesign

Figure 1-3 shows the original curve of our theoretical system, curve A. The dashed line shows the system's original level of perfection. Our system, however, has now been redesigned to a higher level of perfection; that is, a higher level of reliability with a corresponding decrease in total entropy. During this redesign, new components, new materials, or new techniques may have been used to reduce the natural entropy of the system. In some cases, a reduction in

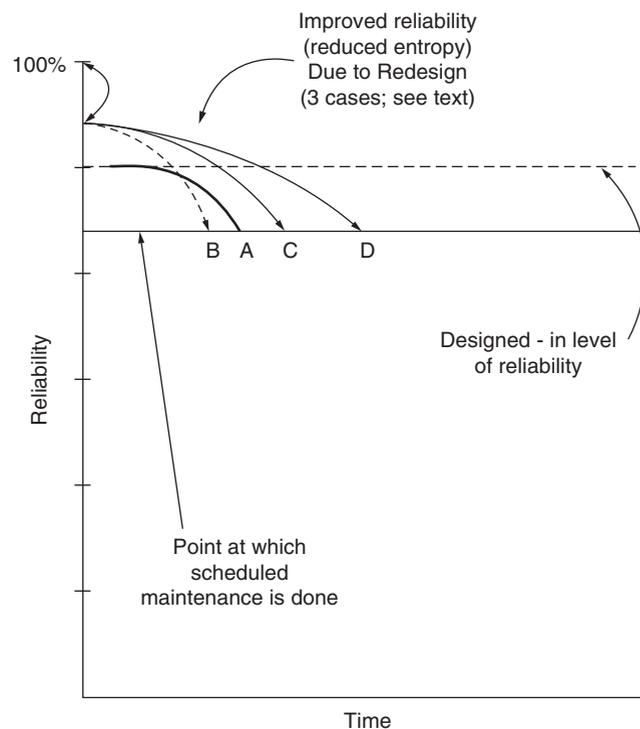


Figure 1-3 Effects of redesign on system reliability.

man-made entropy may result because the designer applied tighter tolerances, attained improved design skills, or changed the design philosophy.

Although the designers have reduced the entropy of the system, the system will still deteriorate. It is quite possible that the rate of deterioration will change from the original design depending upon numerous factors; thus, the slope of the curve may increase, decrease, or stay the same. Whichever is the case, the maintenance requirements of the system could be affected in some way.

If the decay is steeper, as in (B) in Fig. 1-3, the point at which preventive maintenance needs to be performed might occur sooner, and the interval between subsequent actions would be shorter. The result is that maintenance will be needed more often. In this case, the inherent reliability is increased, but more maintenance is required to maintain that level of reliability (level of perfection). Unless the performance characteristics of the system have been improved, this redesign may not be acceptable. A decision must be made to determine if the performance improvement justifies more maintenance and thus an increase in maintenance costs.

Conversely, if the decay rate is the same as before, as shown in curve C of Fig. 1-3, or less steep, as shown in curve D, then the maintenance interval would be increased and the overall amount of preventive maintenance might be reduced. The question to be considered, then, is this: Does the reduction of maintenance justify the cost of the redesign? This question, of course, is a matter for the designers to ponder, not the maintenance people.

One of the major factors in redesign is cost. Figure 1-4 shows the graphs of two familiar and opposing relationships. The upper curve is logarithmic. It represents the increasing perfection attained with more sophisticated design efforts.

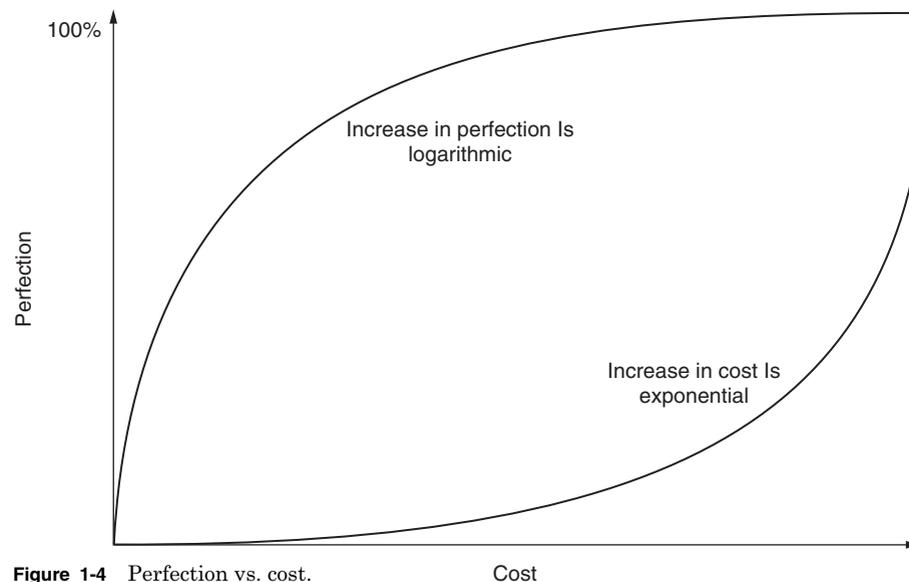
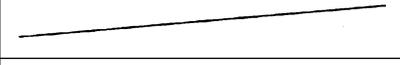
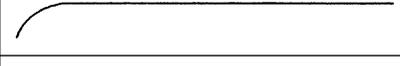
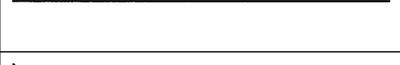
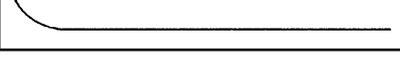


Figure 1-4 Perfection vs. cost.

Cost

10 Fundamentals of Maintenance

TABLE 1-1 Failure Rate Patterns

	A. Infant mortality; constant or slightly rising failure rate; definite wear-out period (4%)
	B. No infant mortality; slightly rising failure rate; definite wear-out period (2%)
	C. No infant mortality; slightly rising rate; no definite wear-out period (5%)
	D. Increasing failure rate at outset; constant or slightly rising rate; no definite wear-out period (7%)
	E. No infant mortality; constant failure rate throughout life; no definite wear-out period (14%)
	F. Infant mortality; constant failure rate throughout life; no definite wear-out period (68%)

Source: F. Stanley Nowlan and Howard F. Heap: *Reliability-Centered Maintenance*; U.S. Department of Commerce, National Technical Information Service, Washington, DC, 1978.

The closer we get to perfection (top of the illustration) the harder it is to make a substantial increase. (We will never get to 100 percent.) The lower curve depicts the cost of those ongoing efforts to improve the system. This, unfortunately, is an exponential curve. The more we try to approach perfection, the more it is going to cost us. It is obvious, then, that the designers are limited in their goal of perfection, not just by entropy but also by costs. The combination of these two limitations is basically responsible for our profession of maintenance.

Failure Rate Patterns

Maintenance, of course, is not as simple as one might conclude from the above discussion of entropy. There is one important fact that must be acknowledged: not all systems or components fail at the same rate nor do they all exhibit the same pattern of wear out and failure. As you might expect, the nature of the maintenance performed on these components and systems is related to those failure rates and failure patterns.

United Airlines did some studies on lifetime failure rates and found six basic patterns.² These are shown in Table 1-1. The vertical axes show failure rates and the horizontal axes indicate time. No values are shown on the scales since these are not really important to the discussion.

Curve A shows what is commonly referred to as the “bathtub” curve, for obvious reasons. This failure rate pattern exhibits a high rate of failure during the early portion of the component’s life, known as infant mortality. This is one of the bugaboos of engineering. Some components exhibit early failures for several

²Nowlan, F. Stanley and Howard F. Heap, *Reliability-Centered Maintenance*. National Technical Information Service, Washington, DC, 1978.

reasons: poor design, improper parts, or incorrect usage. Once the bugs are worked out and the equipment settles into its pattern, the failure rate levels off or rises only slightly over time. That is, until the later stages of the component's life. The rapid rise shown in curve A near the end of its life is an indication of wear out. The physical limit of the component's materials has been reached.

Curve B exhibits no infant mortality but shows a level, or slightly rising failure rate characteristic throughout the component's life until a definite wear-out period is exhibited toward the end.

Curve C depicts components with a slightly increasing failure rate with no infant mortality and no discernible wear-out period, but at some point, it becomes unusable.

Curve D shows a low failure rate when new (or just out of the shop), which rises to some steady level and holds throughout most of the component's life.

Curve E is an ideal component: no infant mortality and no wear-out period, just steady (or slightly rising) failure rate throughout its life.

Curve F shows components with an infant mortality followed by a level or slightly rising failure rate and no wear-out period.

The United Airlines study showed that only about 11 percent of the items included in the experiment (those shown in curves A, B, and C of Table 1-1) would benefit from setting operating limits or from applying a repeated check of wear conditions. The other 89 percent would not. Thus, time of failure or deterioration beyond useful levels could be predicted on only 11 percent of the items (curves A, B, and C of Table 1-1). The other 89 percent (depicted by curves D, E, and F of Table 1-1) would require some other approach. The implication of this variation is that the components with definite life limits and/or wear-out periods will benefit from scheduled maintenance. They will not all come due for maintenance or replacement at the same time, however, but they can be scheduled; and the required maintenance activity can be spread out over the available time, thus avoiding peaks and valleys in the workload. The other 89 percent, unfortunately, will have to be operated to failure before replacement or repair is done. This, being unpredictable, would result in the need for maintenance at odd times and at various intervals; i.e., unscheduled maintenance.

These characteristics of failure make it necessary to approach maintenance in a systematic manner, to reduce peak periods of unscheduled maintenance. The industry has taken this into consideration and has employed several techniques in the design and manufacturing of aircraft and systems to accommodate the problem. These are discussed in the next section.

Other Maintenance Considerations

The aviation industry has developed three management techniques for addressing the in-service interruptions created by the items that must be operated to failure before maintenance can be done. These are equipment redundancy, line replaceable units, and minimum aircraft dispatch requirements.

12 Fundamentals of Maintenance

The concept of redundancy of certain components or systems is quite common in engineering design of systems where a high reliability is desirable. In the case of redundant units—usually called primary and backup units—if one unit fails, the other is available to take over the function. For example, in aviation most commercial jets have two high-frequency (HF) radios. Only one is needed for communications, but the second one is there for backup in case the first one fails.

A unique feature of redundant units also affects the maintenance requirements. If both primary and backup units are instrumented such that the flight crew is aware of any malfunction, no prior maintenance check is required to indicate that incapability. On the other hand, if neither system is so instrumented, maintenance personnel would need to perform some check on both primary and backup systems (at the transit or other check) to determine serviceability.

Very often, however, one system (usually the backup) is instrumented to show serviceability to the crew. If a maintenance check is performed on the other (i.e., the primary) the crew can be assured that it is serviceable. In the case of failure, then, they already have a positive indication, through the instrumentation, that the backup system is available and useable. The purpose for this arrangement is to strike a balance between how much instrumentation is used and how much maintenance is required to ensure system serviceability. In some cases, the backup system is automatically switched into service when the primary system fails. Flight crew needs during the flight are primary concerns in making such decisions.

Another common concept used in aviation is the *line replaceable unit* (LRU). An LRU is a component or system that has been designed in such a manner that the parts that most commonly fail can be quickly removed and replaced on the vehicle. This allows the vehicle to be returned to scheduled service without undue delay for maintenance. The failed part, then, can either be discarded or repaired in the shop as necessary without further delaying the flight.

The third concept for minimizing delays for maintenance in aviation is known as the *minimum equipment list* (MEL). This list allows a vehicle to be dispatched into service with certain items inoperative provided that the loss of function does not affect the safety and operation of the flight. These items are carefully determined by the manufacturer and sanctioned by the regulatory authority during the early stages of vehicle design and test. The manufacturer issues a *master minimum equipment list* (MMEL) which includes all equipment and accessories available for the aircraft model. The airline then tailors the document to its own configuration to produce the MEL (more on this in Chap. 5). Many of these MEL items are associated with redundant systems. The concept of the MEL allows deferral of maintenance without upsetting the mission requirements. The maintenance, however, must be performed within certain prescribed periods, commonly 1, 3, 10, or 30 days, depending on the operational requirements for the system.

The items are identified in the MMEL by flight crew personnel during the latter stages of new aircraft development. Thus, flight personnel determine what systems they can safely fly the mission without or in a degraded condition.

These flight crew personnel also determine how long (1, 3, 10, or 30 days) they can tolerate this condition. Although this is determined in general terms prior to delivering the airplane, the flight crew on board makes the final decision based on actual conditions at the time of dispatch. The pilot in command (PIC) can, based on existing circumstances, decide not to dispatch until repairs are made or can elect to defer maintenance per the airline's MEL. Maintenance must abide by that decision.

Associated with the MEL is a *dispatch deviation guide* (DDG) that contains instructions for the line maintenance crew when the deviation requires some maintenance action that is not necessarily obvious to the mechanic. A dispatch deviation guide is published by the airplane manufacturer to instruct the mechanic on these deviations. The DDG contains such information as tying up cables and capping connectors from removed units, opening and placarding circuit breakers to prevent inadvertent power-up of certain equipment during flight, and any other maintenance action that needs to be taken for precautionary reasons. Similar to the MEL is a *configuration deviation list* (CDL). This list provides information on dispatch of the airplane in the event that certain panels are missing or when other configuration differences not affecting safety are noted. The nonessential equipment and furnishing (NEF) items list contains the most commonly deferred items that do not affect airworthiness or safety of the flight of the aircraft. This is also a part of the MEL system.

Although failures on these complex aircraft can occur at random and can come at inopportune times, these three management actions—redundancy of design, line replaceable units, and minimum dispatch requirements—can help to smooth out the workload and reduce service interruptions.

Establishing a Maintenance Program

Although there has been a considerable amount of improvement in the quality and reliability of components and systems, as well as in materials and procedures, over the 100-year life of aviation, we still have not reached total perfection. Aviation equipment, no matter how good or how reliable, still needs attention from time to time.

Scheduled maintenance and servicing are needed to ensure the designed-in level of perfection (reliability). Due to the nature of the real world, some of these components and systems will, sooner or later, deteriorate beyond a tolerable level or will fail completely. In other instances, users, operators, or even maintenance people who interface with these components and systems can misuse or even abuse the equipment to the extent of damage or deterioration that will require the need for some sort of maintenance action.

We have seen that components and systems fail in different ways and at different rates. This results in a requirement for unscheduled maintenance that is somewhat erratic and uncertain. There are often waves of work and no-work periods that need to be managed to smooth out the workload and stabilize the manpower requirements.

14 Fundamentals of Maintenance

Those components exhibiting life limits or measurable wear-out characteristics can be part of a systematic, scheduled maintenance program. Design redundancy, line replaceable units, and minimum dispatch requirements have been established as management efforts to smooth out maintenance workload. But there are numerous components and systems on an aircraft that do not lend themselves to such adjustment for convenience. Occasionally, inspections and/or modifications of equipment are dictated—within specified time limits—by aviation regulators as well as by manufacturers. It is necessary, then, that the maintenance and engineering organization of an airline be prepared to address the maintenance of aircraft and aircraft systems with a well-thought-out and well-executed program. The remainder of this textbook will address the multifaceted process known as aircraft maintenance and engineering.

The program discussed herein has been created over the years by concentrated and integrated efforts by pilots, airlines, maintenance people, manufacturers, component and system suppliers, regulatory authorities, and professional and business organizations within the aviation industry. Not every airline will need to be organized and operated in the same manner or style, but the programs and activities discussed in this text will apply to all operators.

Chapter
18
Reliability

Introduction

Reliability equals consistency. It can be defined as the probability that an item will perform a required function, under specified conditions without failure, for a specified amount of time according to its intended design. The reliability program is a valuable means of achieving better operational performance in an aircraft maintenance environment, and it is designed to decrease maintenance-related issues and increase flight safety. The intent of this program is to deal systematically with problems as they arise instead of trying to cure immediate symptoms. This program is normally customized, depending on the operators, to accurately reflect the specific operation requirements. Although the word *reliability* has many meanings, in this book we will define the terms that have specialized meanings to aviation maintenance and engineering. In the case of reliability, we first must discuss one important difference in the application of the term.

There are two main approaches to the concept of reliability in the aviation industry. One looks essentially at the whole airline operation or the M&E operation within the whole, and the other looks at the maintenance program in particular. There is nothing wrong with either of these approaches, but they differ somewhat, and that difference must be understood.

The first approach is to look at the overall airline reliability. This is measured essentially by dispatch reliability; that is, by how often the airline achieves an on-time departure¹ of its scheduled flights. Airlines using this approach track delays. Reasons for the delay are categorized as maintenance, flight operations, air traffic control (ATC), etc. and are logged accordingly. The M&E organization is concerned only with those delays caused by maintenance.

¹ On-time departure means that the aircraft has been “pushed back” from the gate within 15 minutes of the scheduled departure time.

218 Oversight Functions

Very often, airlines using this approach to reliability overlook any maintenance problems (personnel or equipment related) that do not cause delays, and they track and investigate only those problems that do cause delays. This is only partially effective in establishing a good maintenance program.

The second approach (which we should actually call the primary approach) is to consider reliability as a program specifically designed to address the problems of maintenance—whether or not they cause delays—and provide analysis of and corrective actions for those items to improve the overall reliability of the equipment. This contributes to the dispatch reliability, as well as to the overall operation.

We are not going to overlook the dispatch reliability, however. This is a distinct part of the reliability program we discuss in the following pages. But we must make the distinction and understand the difference. We must also realize that not all delays are caused by maintenance or equipment even though maintenance is the center of attention during such a delay. Nor can we only investigate equipment, maintenance procedures, or personnel for those discrepancies that have caused a delay. As you will see through later discussions, dispatch reliability is a subset of overall reliability.

Types of Reliability

The term *reliability* can be used in various respects. You can talk about the overall reliability of an airline's activity, the reliability of a component or system, or even the reliability of a process, function, or person. Here, however, we will discuss reliability in reference to the maintenance program specifically.

There are four types of reliability one can talk about related to the maintenance activity. They are (a) statistical reliability, (b) historical reliability, (c) event-oriented reliability, and (d) dispatch reliability. Although dispatch reliability is a special case of event-oriented reliability, we will discuss it separately due to its significance.

Statistical reliability

Statistical reliability is based upon collection and analysis of failure, removal, and repair rates of systems or components. From this point on, we will refer to these various types of maintenance actions as "events." Event rates are calculated on the basis of events per 1000 flight hours or events per 100 flight cycles. This normalizes the parameter for the purpose of analysis. Other rates may be used as appropriate.

Many airlines use statistical analysis, but some often give the statistics more credence than they deserve. For one example, airlines with 10 or more aircraft tend to use the statistical approach, but most teachers and books on statistics tell us that for any data set with less than about 30 data points the statistical calculations are not very significant. Another case of improper use of statistics was given as an example presented in an aviation industry seminar on reliability.

The airline representative used this as an example of why his airline was going to stop using statistical reliability. Here is his example.

We use weather radar only two months of the year. When we calculate the mean value of failure rates and the alert level in the conventional manner [discussed in detail later in this chapter] we find that we are always on alert. This, of course, is not true.

The gentleman was correct in defining an error in this method, and he was correct in determining that—at least in this one case—statistics was not a valid approach. Figure 18-1 shows why.

The top curve in Fig. 18-1 shows the two data points for data collected when the equipment was in service. It also shows 10 zero data points for those months when the equipment was not used and no data were collected (12-month column). These zeros are not valid statistical data points. They do not represent zero failures; they represent “no data” and therefore should not be used in the calculation. Using these data, however, has generated a mean value (lower, dashed line) of 4.8 and an alert level at two standard deviations above the mean (upper, solid line) of 27.6.

One thing to understand about mathematics is that the formulas will work, will produce numerical answers, whether or not the input data are correct. Garbage in, garbage out. The point is, you only have two valid data points here shown in the bottom curve of Fig. 18-1 (2-month data). The only meaningful statistic here is the average of the two numbers, 29 (dashed line). One can calculate a standard deviation (SD) here using the appropriate formula or a calculator, but the parameter has no meaning for just two data points. The alert level set

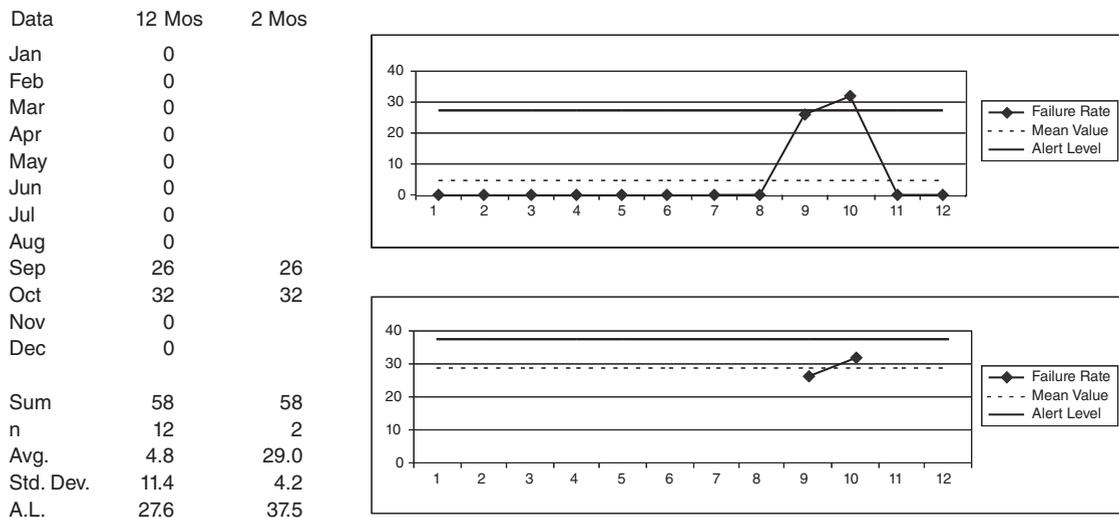


Figure 18-1 Comparison of alert level calculation methods.

220 Oversight Functions

by using this calculation is 37.5 (solid line). For this particular example, statistical reliability is not useable, but historical reliability is quite useful. We will discuss that subject in the next section.

Historical reliability

Historical reliability is simply a comparison of current event rates with those of past experience. In the example of Fig. 18-1, the data collected show fleet failures of 26 and 32 for the 2 months the equipment was in service. Is that good or bad? Statistics will not tell you but history will. Look at last year's data for the same equipment, same time period. Use the previous year's data also, if available. If current rates compare favorably with past experience, then everything is okay; if there is a significant difference in the data from one year to the next, that would be an indication of a possible problem. That is what a reliability program is all about: detecting and subsequently resolving problems.

Historical reliability can be used in other instances, also. The most common one is when new equipment is being introduced (components, systems, engines, aircraft) and there is no previous data available on event rates, no information on what sort of rates to expect. What is "normal" and what constitutes "a problem" for this equipment? In historical reliability we merely collect the appropriate data and literally "watch what happens." When sufficient data are collected to determine the "norms," the equipment can be added to the statistical reliability program.

Historical reliability can also be used by airlines wishing to establish a statistically based program. Data on event rates kept for 2 or 3 years can be tallied or plotted graphically and analyzed to determine what the normal or acceptable rates would be (assuming no significant problems were incurred). Guidelines can then be established for use during the next year. This will be covered in more detail in the reliability program section below.

Event-oriented reliability

Event-oriented reliability is concerned with one-time events such as bird strikes, hard landings, overweight landings, in-flight engine shutdowns, lighting strikes, ground or flight interruption, and other accidents or incidents. These are events that do not occur on a daily basis in airline operations and, therefore, produce no usable statistical or historical data. Nevertheless, they do occur from time to time, and each occurrence must be investigated to determine the cause and to prevent or reduce the possibility of recurrence of the problem.

In ETOPS² operations, certain events associated with this program differ from conventional reliability programs, and they do rely on historical data and alert levels to determine if an investigation is necessary to establish whether a problem can be reduced or eliminated by changing the maintenance program.

² Requirements for extended range operations with two-engine airplanes (ETOPS) are outlined in FAA Advisory Circular AC 120-42B, and also discussed in Appendix E of this book.

Events that are related to ETOPS flights are designated by the FAA as actions to be tracked by an “event-oriented reliability program” in addition to any statistical or historical reliability program. Not all the events are investigated, but everything is continually monitored in case a problem arises.

Dispatch reliability

Dispatch reliability is a measure of the overall effectiveness of the airline operation with respect to on-time departure. It receives considerable attention from regulatory authorities, as well as from airlines and passengers, but it is really just a special form of the event-oriented reliability approach. It is a simple calculation based on 100 flights. This makes it convenient to relate dispatch rate in percent. An example of the dispatch rate calculation follows.

If eight delays and cancellations are experienced in 200 flights, that would mean that there were four delays per 100 flights, or a 4 percent delay rate. A 4 percent delay rate would translate to a 96 percent dispatch rate (100 percent – 4 percent delayed = 96 percent dispatched on time). In other words, the airline dispatched 96 percent of its flights on time.

The use of dispatch reliability at the airlines is, at times, misinterpreted. The passengers are concerned with timely dispatch for obvious reasons. To respond to FAA pressures on dispatch rate, airlines often overreact. Some airline maintenance reliability programs track only dispatch reliability; that is, they only track and investigate problems that resulted in a delay or a cancellation of a flight. But this is only part of an effective program and dispatch reliability involves more than just maintenance. An example will bear this out.

The aircraft pilot in command is 2 hours from his arrival station when he experiences a problem with the rudder controls. He writes up the problem in the aircraft logbook and reports it by radio to the flight following unit at the base. Upon arrival at the base, the maintenance crew meets the plane and checks the log for discrepancies. They find the rudder control write-up and begin troubleshooting and repair actions. The repair takes a little longer than the scheduled turnaround time and, therefore, causes a delay. Since maintenance is at work and the rudder is the problem, the delay is charged to maintenance and the rudder system would be investigated for the cause of the delay.

This is an improper response. Did maintenance cause the delay? Did the rudder equipment cause the delay? Or was the delay caused by poor airline procedures? To put it another way: could a change of airline procedures eliminate the delay? Let us consider the events as they happened and how we might change them for the better.

If the pilot and the flight operations organization knew about the problem 2 hours before landing, why wasn't maintenance informed at the same time? If they had been informed, they could have spent the time prior to landing in studying the problem and performing some troubleshooting analysis. It is quite possible, then, that when the airplane landed, maintenance could have met it with a fix in hand. Thus, this delay could have been prevented by procedural changes. The procedure should be changed to avoid such delays in the future.

222 Oversight Functions

While the maintenance organization and the airline could benefit from this advance warning of problems, it will not always eliminate delays. The important thing to remember is that if a delay is caused by procedure, it should be attributed to procedure and it should be avoided in the future by altering the procedure. That is what a reliability program is about: detecting where the problems are and correcting them, regardless of who or what is to blame.

Another fallacy in overemphasizing dispatch delay is that some airlines will investigate each delay (as they should), but if an equipment problem is involved, the investigation may or may not take into account other similar failures that did not cause delays. For example, if you had 12 write-ups of rudder problems during the month and only one of these caused a delay, you actually have two problems to investigate: (a) the delay, which could be caused by problems other than the rudder equipment and (b) the 12 rudder write-ups that may, in fact, be related to an underlying maintenance problem. One must understand that dispatch delay constitutes one problem and the rudder system malfunction constitutes another. They may indeed overlap but they are two different problems. The delay is an event-oriented reliability problem that must be investigated on its own; the 12 rudder problems (if this constitutes a high failure rate) should be addressed by the statistical (or historical) reliability program. The investigation of the dispatch delays should look at the whole operation. Equipment problems—whether or not they caused delays—should be investigated separately.

A Reliability Program

A reliability program for our purposes is, essentially, a set of rules and practices for managing and controlling a maintenance program. The main function of a reliability program is to monitor the performance of the vehicles and their associated equipment and call attention to any need for corrective action. The program has two additional functions: (a) to monitor the effectiveness of those corrective actions and (b) to provide data to justify adjusting the maintenance intervals or maintenance program procedures whenever those actions are appropriate.

Elements of a Reliability Program

A good reliability program consists of seven basic elements as well as a number of procedures and administrative functions. The basic elements (discussed in detail below) are (a) data collection; (b) problem area alerting, (c) data display; (d) data analysis; (e) corrective actions; (f) follow-up analysis; and (g) a monthly report. We will look at each of these seven program elements in more detail.

Data collection

We will list 10 data types that can be collected, although they may not necessarily be collected by all airlines. Other items may be added at the airline's discretion.

The data collection process gives the reliability department the information needed to observe the effectiveness of the maintenance program. Those items that are doing well might be eliminated from the program simply because the data show that there are no problems. On the other hand, items not being tracked may need to be added to the program because there are serious problems related to those systems. Basically, you collect the data needed to stay on top of your operation. The data types normally collected are as follows:

1. Flight time and cycles for each aircraft
2. Cancellations and delays over 15 minutes
3. Unscheduled component removals
4. Unscheduled engine removals
5. In-flight shutdowns of engines
6. Pilot reports or logbook write-ups
7. Cabin logbook write-ups
8. Component failures (shop maintenance)
9. Maintenance check package findings
10. Critical failures

We will discuss each of these in detail below.

Flight time and flight cycles. Most reliability calculations are “rates” and are based on flight hours or flight cycles; e.g., 0.76 failures per 1000 flight hours or 0.15 removals per 100 flight cycles.

Cancellations and delays over 15 minutes. Some operators collect data on all such events, but maintenance is concerned primarily with those that are maintenance related. The 15-minute time frame is used because that amount of time can usually be made up in flight. Longer delays may cause schedule interruptions or missed connections, thus the need for rebookings. This parameter is usually converted to a “dispatch rate” for the airline as discussed above.

Unscheduled component removals. This is the unscheduled maintenance mentioned earlier and is definitely a concern of the reliability program. The rate at which aircraft components are removed may vary widely depending on the equipment or system involved. If the rate is not acceptable, an investigation should be made and some sort of corrective action must be taken. Components that are removed and replaced on schedule—e.g., HT items and certain OC items—are not included here, but these data may be collected to aid in justifying a change in the HT or OC interval schedule.

Unscheduled removals of engines. This is the same as component removals, but obviously an engine removal constitutes a considerable amount of time and manpower; therefore, these data are tallied separately.

In-flight shutdown (IFSD) of engines. This malfunction is probably one of the most serious in aviation, particularly if the airplane only has two engines (or one).

224 Oversight Functions

The FAA requires a report of IFSD within 72 hours.³ The report must include the cause and the corrective action. The ETOPS operators are required to track IFSDs and respond to excessive rates as part of their authorization to fly ETOPS. However, non-ETOPS operators also have to report shutdowns and should also be tracking and responding to high rates through the reliability program.

Pilot reports or logbook write-ups. These are malfunctions or degradations in airplane systems noted by the flight crew during flight. Tracking is usually by ATA Chapter numbers using two, four, or six digits. This allows pinpointing of the problems to the system, subsystem, or component level as desired. Experience will dictate what levels to track for specific equipment.

Cabin logbook write-ups. These discrepancies may not be as serious as those the flight crew deals with, but passenger comfort and the ability of the cabin crew to perform their duties may be affected. These items may include cabin safety inspection, operational check of cabin emergency lights, first aid kits, and fire extinguishers. If any abnormality is found, these items are written up by the flight crew in the maintenance logbook as a discrepancy item.

Component failures. Any problems found during shop maintenance visits are tallied for the reliability program. This refers to major components within the black boxes (avionics) or parts and components within mechanical systems.

Maintenance check package findings. Systems or components found to be in need of repair or adjustment during normal scheduled maintenance checks (non-routine items) are tracked by the reliability program.

Critical failures. Failures involving a loss of function or secondary damage that could have a direct adverse effect on operating safety.

Problem detection—an alerting system

The data collection system allows the operator to compare present performance with past performance in order to judge the effectiveness of maintenance and the maintenance program. An alerting system should be in place to quickly identify those areas where the performance is significantly different from normal. These are items that might need to be investigated for possible problems. Standards for event rates are set according to analysis of past performances and deviations from these standards.

This alert level is based on a statistical analysis of the event rates of the previous year, offset by 3 months. The mean value of the failure rates and the standard deviation from the mean are calculated, and an alert level is set at one to three standard deviations above that mean rate (more on setting and adjusting alert levels later). This value, the upper control limit (UCL), is commonly referred to as the alert level. However, there is an additional calculation that can be made to smooth the curve and help eliminate “false alerts.” This is the 3-month rolling average, or trend line. The position of these two lines (the monthly rate and the 3-month average) relative to the UCL is used to determine alert status.

³ See Federal Aviation Regulation 121.703, Mechanical Reliability Report.

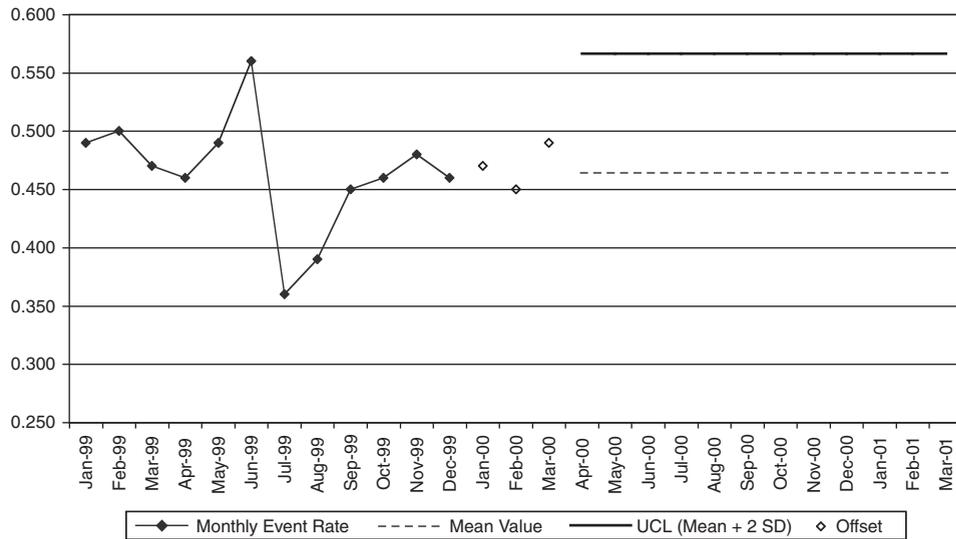


Figure 18-2 Calculation of new alert levels.

Setting and adjusting alert levels

It is recommended that alert levels be recalculated yearly. The data used to determine alert level are the event rates for the previous year offset by 3 months. The reason for this will be explained shortly.

Figure 18-2 shows the data used and the results in graphic form. In this example, we are establishing a new alert level for the year April 2000 through March 2001. This level is represented in Fig. 18-2 as the upper straight line. These data were obtained using the actual event rates for January 1999 through December 1999 shown on the left of the figure. The three data points between (shown as diamonds for January to March 2000 in Fig. 18-2) will be used in calculating a 3-month rolling average to be used during the collection of new data. This will be discussed later.

Basic statistics are used for the calculations. From the original data (January–December 1999) we calculate the mean and the standard deviation of these data points. The mean is used as a baseline for the new data and is shown as the dashed line on the right side of Fig. 18-2. The solid line on the right of Fig. 18-2 is the alert level that we have chosen for these data and is equal to the calculated mean plus two standard deviations. Event rates for the new year, then, will be plotted and measured relative to these guidelines.

Reading alert status

The data shown in Fig. 18-3 show 1 year of event rates (solid jagged line with triangles) along with the mean value (bottom straight line) and the alert level (upper straight line). As you can see, the event rate swings above the alert level several times through the year (February, June, October, and December). Of course,

226 Oversight Functions

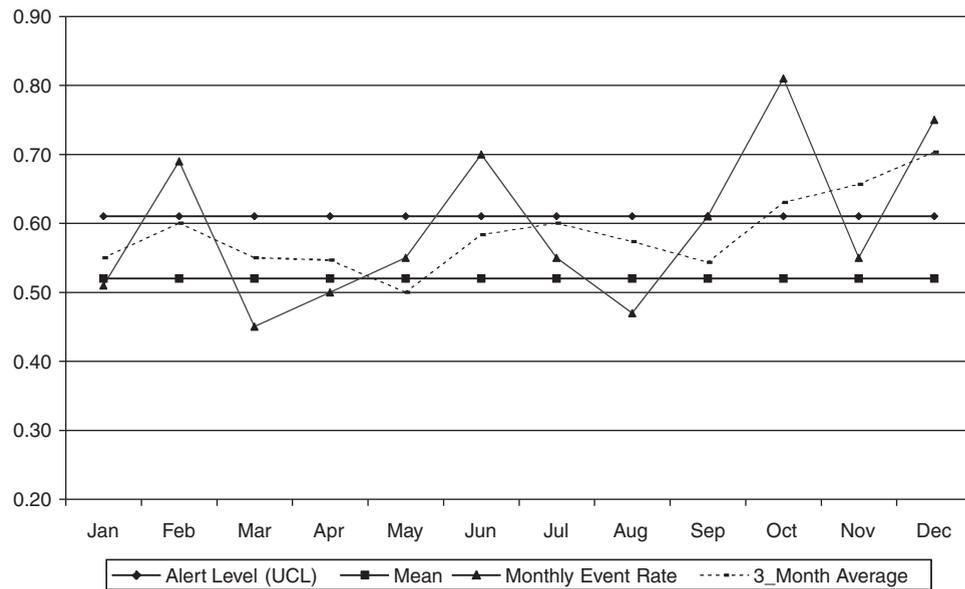


Figure 18-3 Reading alert status.

it is easy to see the pattern as we look at the year's events. But in reality, you will only see 1 month at a time and the preceding months. Information on what is going to happen the next month is not available to you.

When the event rate goes above the alert level (as in February), it is not necessarily a serious matter. But if the rate stays above the alert level for 2 months in succession, then it may warrant an investigation. The preliminary investigation may indicate a seasonal variation or some other one-time cause, or it may suggest the need for a more detailed investigation. More often than not, it can be taken for what it was intended to be—an “alert” to a possible problem. The response would be to wait and see what happens next month. In Fig. 18-3, the data show that, in the following month (March) the rate went below the line; thus, no real problem exists. In other words, when the event rate penetrates the alert level, it is not an indication of a problem; it is merely an “alert” to the possibility of a problem. Reacting too quickly usually results in unnecessary time and effort spent in investigation. This is what we call a “false alert.”

If experience shows that the event rate for a given item varies widely from month to month above and below the UCL as in Fig. 18-3—and this is common for some equipment—many operators use a 3-month rolling average. This is shown as the dashed line in Fig. 18-3. For the first month of the new data year, the 3-month average is determined by using the offset data points in Fig. 18-2. (Actually, only 2 months offset is needed, but we like to keep things on a quarterly basis.) The purpose for the offset is to ensure that the plotted data for the new year do not contain any data points that were used to determine the mean and alert levels we use for comparison.

While the event rate swings above and below the alert level, the 3-month rolling average (dashed line) stays below it—until October. This condition—event rate and 3-month average above the UCL—indicates a need to watch the activity more closely. In this example, the event rate went back down below the UCL in November, but the 3-month average stayed above the alert level. This is an indication that the problem should be investigated.

Setting alert levels

These upper control limits, or alert levels, and the mathematics that produced them are not magical by any means. They will not tell you when you have a definite problem nor will they tell you where or what to investigate. What they will do is provide you with intelligent guidelines for making your own decisions about how to proceed. The whole process begins with your intellect and your ability to set these alert levels to an effective level.

Earlier in this chapter, we talked about an airline that was rejecting statistical reliability and gave an example of why. Another of the reasons the gentleman gave for this decision was that “we know we have problems with engines, but engines are never on alert.” If you use the UCL concept to alert you to possible problems and you do not get an alert indication when you know you have problems, then it should not take much thought to make you realize that your chosen alert level is wrong. This alert level is a very important parameter and it must be set to a useable level, a level that will indicate to you that a problem exists or may be developing. If not set properly, the alert level is useless. And that is not the fault of statistics.

This use of an alert level is designed to tell you when you have (or may have) a problem developing that requires investigation. But you have to know what conditions constitute a possible problem and set the alert level accordingly. You have to know your equipment and its failure patterns to determine when you should proceed with an investigation and when to refrain from investigating. You have to recognize “false alerts.” You also have to know whether or not the event rate data points for a particular item are widely or narrowly distributed; that is, if it has a large or small standard deviation. This knowledge is vital to setting useable alert levels.

Many airlines erroneously set all alert levels at two standard deviations above the mean. Unfortunately, this is not a good practice. It is a good place to start, but there must be an adjustment in some cases to provide the most useable data and to avoid false alerts.

As we discussed in Chap. 1, not everything fails at the same rate or in the same pattern. Event rates tracked by a reliability program can be quite erratic, as the data in Fig. 18-3 show. For other rates, the numbers can be more stable. This characteristic of the data is depicted by the statistical parameter of standard deviation—the measure of the distribution of data points around the mean. A large standard deviation means wide distribution, a large variation in point values. A small standard deviation means that the points are closer together.

228 Oversight Functions

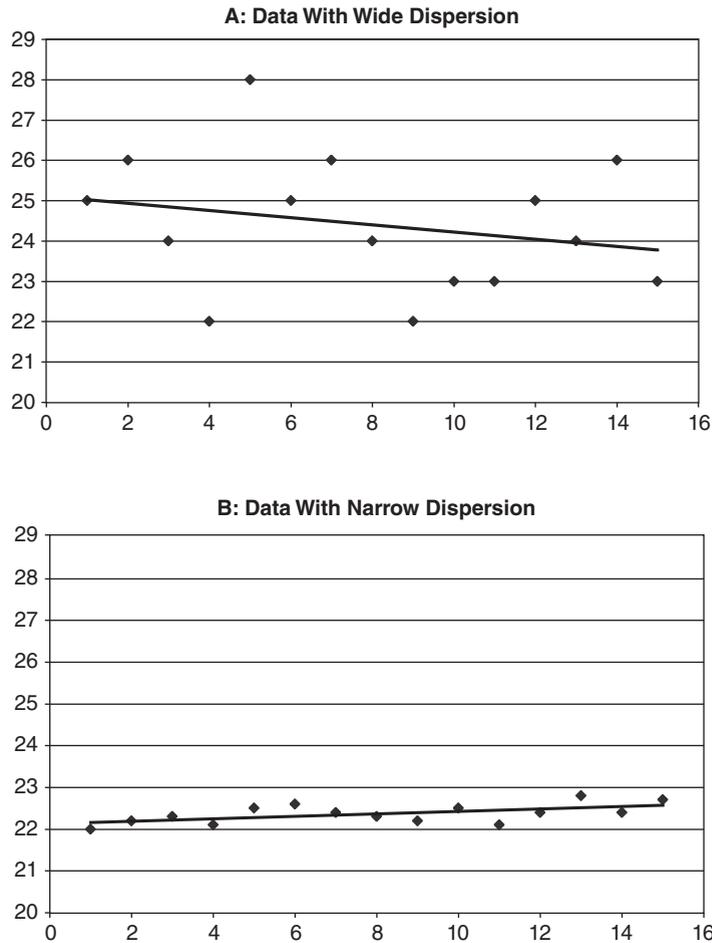


Figure 18-4 Dispersion of data points.

Figure 18-4 shows the difference between two data sets. The data points in (A) are widely scattered or distributed about the mean while those in (B) are all very close together around the mean. Note that the averages of these two data sets are nearly equal but the standard deviations are quite different. Figure 18-5 shows the bell-shaped distribution curve. One, two, and three standard deviations in each case are shown on the graph. You can see here that, at one SD only 68 percent of the valid failure rates are included. At two standard deviations above the mean, you still have not included all the points in the distribution. In fact, two standard deviations above and below the mean encompass only 95.5 percent of the points under the curve; i.e., just over 95 percent of the valid failure rates. This is why we do not consider an event rate in this range a definite problem. If it remains above this level in the following month it may suggest a possible problem.

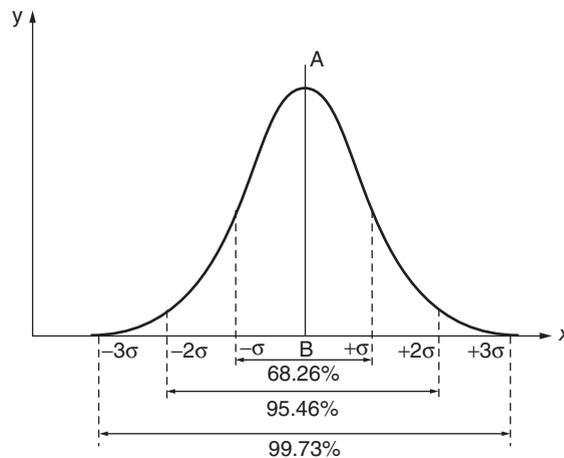


Figure 18-5 Standard bell-shaped curve. (Source: *The Standard Handbook for Aeronautical and Astronautical Engineers*, New York, NY: McGraw-Hill, 2003.)

On the other hand, if the event rate data you are working with had a small standard deviation, it would be difficult to distinguish between two and three SDs. In this case, the alert level should be set at three SDs.

This alert level system can be overdone at times. The statistics used are not exact. We are assuming that the event rates will always have a distribution depicted by the bell-shaped curve. We assume that our data are always accurate and that our calculations are always correct. But this may not be true. These alert levels are merely guidelines to identifying what should be investigated and what can be tolerated. Use of the alert level is not rocket science but it helps ease the workload in organizations with large fleets and small reliability staffs. Some airlines, using only event rates, will investigate perhaps the 10 highest rates; but this does not always include the most important or the most significant equipment problems. The alert level approach allows you to prioritize these problems and work the most important ones first.

Data display

Several methods for displaying data are utilized by the reliability department to study and analyze the data they collect. Most operators have personal computers available so that data can easily be displayed in tabular and graphical forms. The data are presented as events per 100 or 1000 flight hours or flight cycles. Some, such as delays and cancellations, are presented as events per 100 departures. The value of 100 allows easy translation of the rate into a percentage.

Tabular data allow the operator to compare event rates with other data on the same sheet. It also allows the comparison of quarterly or yearly data (see Table 18-1). Graphs, on the other hand, allow the operator to view the month-to-month performance and note, more readily, those items that show increasing

230 Oversight Functions

TABLE 18-1 Pilot Reports per 100 Landings (by ATA Chapter)

ATA Chapter	System	PIREPS	June-99	July-99	August-99	Three-month average	UCL	Mean	Alert status
21	Air conditioning	114	3.65	3.77	3.80	3.74	3.75	2.70	YE
22	Auto flight	43	1.80	1.48	1.45	1.58	1.39	1.21	WA
23	Communications	69	3.44	2.75	2.33	2.84	2.80	2.30	CL
24	Electrical power	29	1.15	0.87	0.98	1.00	0.94	0.60	AL
25	Equip/furnishings	104	4.17	3.69	3.52	3.79	5.43	4.38	
26	Fire protection	30	1.80	1.30	1.01	1.37	2.19	1.14	
27	Flight controls	48	0.99	3.07	1.62	1.89	1.94	1.26	
28	Fuel	36	0.65	1.16	1.22	1.01	2.32	1.27	
29	Hydraulic power	17	0.73	0.43	0.57	0.58	1.58	0.82	
30	Ice & rain protection	12	0.61	0.65	0.41	0.56	0.72	0.56	
31	Instruments	49	1.76	1.48	1.66	1.63	2.46	1.66	
32	Landing gear	67	2.41	2.06	2.27	2.25	2.72	1.76	
33	Lights	72	3.48	3.15	2.43	3.02	3.32	2.42	
34	Navigation	114	4.81	6.62	3.85	5.09	5.58	4.70	
35	Oxygen	19	0.31	0.67	0.64	0.54	0.41	0.23	YE
36	Pneumatics	25	1.11	0.80	0.85	0.92	1.19	0.77	
38	Water & waste	16	0.42	0.36	0.54	0.44	1.10	0.56	
49	Aux. power	42	1.41	1.48	1.42	1.44	1.63	1.38	
51	Structures	0	0.00	0.00	0.00	0.00	0.16	0.09	
52	Doors	31	1.41	1.05	1.05	1.17	1.62	0.92	
53	Fuselage	0	0.00	0.00	0.00	0.00	0.33	0.02	
54	Nacelles & pylons	1	0.00	0.00	0.08	0.03	0.22	0.10	
55	Stabilizers	0	0.00	0.00	0.00	0.00	0.16	0.09	
56	Windows	0	0.00	0.04	0.00	0.01	0.09	0.06	
57	Wings	0	0.00	0.00	0.00	0.00	0.33	0.15	
71	Power plant	11	0.65	0.54	0.37	0.52	1.30	0.91	
72	Engine	4	0.31	0.29	0.14	0.25	0.47	0.22	
73	Fuel & controls	17	0.96	0.47	0.57	0.67	0.84	0.61	
74	Ignition	11	0.08	0.40	0.37	0.28	0.46	0.30	
75	Air	53	1.52	1.63	1.79	1.65	1.11	0.66	RA
76	Engine control	3	0.23	0.14	0.10	0.16	0.33	0.15	
77	Engine indicating	22	0.53	0.76	0.74	0.68	0.96	0.68	
78	Exhaust	3	0.50	0.43	0.10	0.34	0.90	0.64	
79	Oil	5	0.19	0.22	0.17	0.19	0.83	0.48	
80	Starting	3	0.27	0.29	0.10	0.22	0.28	0.17	CL
	Total	1070							

NOTE: Alert status codes: CL = clear from alert; YE = yellow alert; AL = red alert; RA = remains in alert; WA = watch.

rates and appear to be heading for alert status (see Fig. 18-3). This is a great help in analysis. Some of the data collected may be compared on a monthly basis, by event, or by sampling.

Table 18-1 is a listing of pilot reports (PIREPS) or maintenance logbook entries recorded by a typical airline for 1 month of operation for a fleet of aircraft. The numbers are examples only and do not represent any particular operator, aircraft, or fleet size. For these data, a tally is kept by ATA Chapter, and event rates are calculated as PIREPS per 100 landings. The chart shows data for the current month (August 99) and the two previous months along with the

3-month rolling average. The alert level or UCL and the mean value of event rate, calculated as discussed in the text, are also included. Seven of these ATA Chapters have alert indications noted in the last column.

Chapter 21 has had an event rate above the UCL for 2 months running (July, August); therefore, this represents a yellow alert (YE). Depending on the severity of the problem, this may or may not require an immediate investigation. Chapter 24, however, is different. For July, the event rate was high, 1.15. If this were the first time for such a rate, it would have been listed in the report for that month as a watch (WA). The rate went down in July but has gone up again in August. In the current report, then, it is a full alert condition. It is not only above the alert level, it has been above 2 of the 3 months, and it appears somewhat erratic. It is left as an exercise for the student to analyze the other alert status items. What about ATA Chapter 38?

Data analysis

Whenever an item goes into alert status, the reliability department does a preliminary analysis to determine if the alert is valid. If it is valid, a notice of the on-alert condition is sent to engineering for a more detailed analysis. The engineering department is made up of experienced people who know maintenance and engineering. Their job relative to these alerts is to troubleshoot the problem, determine the required action that will correct the problem, and issue an engineering order (EO) or other official paperwork that will put this solution in place.

At first, this may seem like a job for maintenance. After all, troubleshooting and corrective action is their job. But we must stick with our basic philosophy from Chap.7 of separating the inspectors from the inspected. Engineering can provide an analysis of the problem that is free from any unit bias and be free to look at all possibilities. A unit looking into its own processes, procedures, and personnel may not be so objective. The engineering department should provide analysis and corrective action recommendations to the airline Maintenance Program Review Board (discussed later) for approval and initiation.

Note: Appendix C discusses the troubleshooting process that applies to engineers as well as mechanics; and Appendix D outlines additional procedures for reliability and engineering alert analysis efforts.

Corrective action

Corrective actions can vary from one-time efforts correcting a deficiency in a procedure to the retraining of mechanics to changes in the basic maintenance program. The investigation of these alert conditions commonly results in one or more of the following actions: (a) modifications of equipment; (b) change in or correction to line, hangar, or shop processes or practices; (c) disposal of defective parts (or their suppliers); (d) training of mechanics (refresher or upgrade); (e) addition of maintenance tasks to the program; or (f) decreases in maintenance

232 Oversight Functions

intervals for certain tasks. Engineering then produces an engineering order for implementation of whatever action is applicable. Engineering also tracks the progress of the order and offers assistance as needed. Completion of the corrective action is noted in the monthly reliability report (discussed later). Continual monitoring by reliability determines the effectiveness of the selected corrective action.

Corrective actions should be completed within 1 month of issuance of the EO. Completion may be deferred if circumstances warrant, but action should be completed as soon as possible to make the program effective. Normally, the Maintenance Program Review Board (MPRB) will require justification in writing for extensions of this period; the deferral, and the reason for deferral, will be noted in the monthly report.

Follow-up analysis

The reliability department should follow up on all actions taken relative to on-alert items to verify that the corrective action taken was indeed effective. This should be reflected in decreased event rates. If the event rate does not improve after action has been taken, the alert is reissued and the investigation and corrective action process is repeated, with engineering taking a different approach to the problem. If the corrective action involves lengthy modifications to numerous vehicles, the reduction in the event rate may not be noticeable for some time. In these cases, it is important to continue monitoring the progress of the corrective action in the monthly report along with the ongoing event rate until corrective action is completed on all vehicles. Then follow-up observation is employed to judge the effectiveness (wisdom) of the action. If no significant change is noted in the rates within a reasonable time after a portion of the fleet has been completed, the problem and the corrective action should be reanalyzed.

Data reporting

A reliability report is issued monthly. Some organizations issue quarterly and yearly reports in summary format. The most useful report, however, is the monthly. This report should not contain an excessive amount of data and graphs without a good explanation of what this information means to the airline and to the reader of the report. The report should concentrate on the items that have just gone on alert, those items under investigation, and those items that are in or have completed the corrective action process. The progress of any items that are still being analyzed or implemented will also be noted in the report, showing status of the action and percent of fleet completed if applicable. These items should remain in the monthly report until all action has been completed and the reliability data show positive results.

Other information, such as a list of alert levels (by ATA Chapter or by item) and general information on fleet reliability will also be included in the monthly report. Such items as dispatch rates, reasons for delays and/or cancellations,

flight hours and cycles flown and any significant changes in the operation that affect the maintenance activity would also be included. The report should be organized by fleet; that is, each airplane model would be addressed in a separate section of the report.

The monthly reliability report is not just a collection of graphs, tables, and numbers designed to dazzle higher-level management. Nor is it a document left on the doorstep of others, such as QA or the FAA, to see if they can detect any problems you might have. This monthly report is a working tool for maintenance management. Besides providing operating statistics, such as the number of aircraft in operation, the number of hours flown, and so forth, it also provides management with a picture of what problems are encountered (if any) and what is being done about those problems. It also tracks the progress and effectiveness of the corrective action. The responsibility for writing the report rests with the reliability department, not engineering.

Other Functions of the Reliability Program

Investigation of the alert items by engineering often results in the need to change the maintenance program. This can mean (a) changes in specific tasks; (b) adjustments in the interval at which maintenance tasks are performed; or (c) changes in the maintenance processes (HT, OC, and CM) to which components are assigned. A change in the task may mean rewriting maintenance and/or test procedures or in implementing new, more effective procedures.

Adjustments in the maintenance interval may be a solution to a given problem. A maintenance action currently performed at, say a monthly interval, should, in fact, be done weekly or even daily to reduce the event rate. The reliability program should provide the rules and processes used to adjust these intervals. The Maintenance Program Review Board must approve these changes and, in certain instances, the regulatory authority must also approve. Generally, though, the change to a greater frequency (shorter interval) is not difficult. One should keep in mind, however, that this means higher cost of maintenance due to the increase in maintenance activity. This cost must be offset by the reduction in the event rate that generated the change and a reduction in the maintenance requirements resulting from the change. The economics of this change is one of the concerns engineering must address during the investigation of the alert condition. The cost of the change may or may not be offset by the gain in reliability or performance (see objective 5 in Chap. 3).

Administration and Management of the Reliability Program

On the administration and management side, a reliability program will include written procedures for changing maintenance program tasks, as well as processes and procedures for changing maintenance intervals (increasing or decreasing them). Identification, calculation, establishment, and adjustment of alert levels

234 Oversight Functions

and the determination of what data to track are basic functions of the reliability section. Collecting data is the responsibility of various M&E organizations, such as line maintenance (flight hours and cycles, logbook reports, etc.); overhaul shops (component removals); hangar (check packages); and materiel (parts usage). Some airlines use a central data collection unit for this, located in M&E administration, or some other unit such as engineering or reliability. Other airlines have provisions for the source units to provide data to the reliability department on paper or through the airline computer system. In either case, reliability is responsible for collecting, collating, and displaying these data and performing the preliminary analysis to determine alert status.

The reliability department analyst in conjunction with MCC keeps a watchful eye on the aircraft fleet and its systems for any repeat maintenance discrepancies. The analyst reviews reliability reports and items on a daily basis, including aircraft daily maintenance, time-deferred maintenance items, MEL, and other out of service events with any type of repeat mechanical discrepancies.

The analyst plans a sequence of repair procedures if aircraft have repeated the maintenance discrepancy three times or more and have exhausted any type of fix to rid the aircraft of the maintenance discrepancy. The analyst is normally in contact with the MCC and local aircraft maintenance management to coordinate a plan of attack with the aircraft manufacturer's maintenance help desk to ensure proper tracking and documenting of the actual maintenance discrepancy and corrective action planned or maintenance performed. These types of communications are needed for airlines to run a successful maintenance operation and to keep the aircraft maintenance downtime to a minimum. This normally occurs when a new type of aircraft is added to the airline's fleet. Sometimes maintenance needs help fixing a recurring problem.

Maintenance program review board

The solution of reliability problems is not the exclusive domain of the reliability section or the engineering section; it is a maintenance and engineering organization-wide function. This group approach ensures that all aspects of the problem have been addressed by those who are most familiar with the situation. Therefore, oversight of the program is assigned to a Maintenance Program Review Board that is made up of key personnel in M&E. Based on the typical organization of Chap. 7, the MPRB would consist of the following personnel:

1. Director of MPE as chairman
2. Permanent members
 - a. Director of technical services
 - b. Director of airplane maintenance
 - c. Director of overhaul shops
 - d. Director of QA and QC
 - e. Manager of QA and QC

- f.* Manager of engineering
- g.* Manager of reliability
- 3. Adjunct members are representatives of affected M&E departments
 - a.* Engineering supervisors (by ATA Chapter or specialty)
 - b.* Airplane maintenance (line, hangar)
 - c.* Overhaul shops (avionics, hydraulics, etc.)
 - d.* Production planning and control
 - e.* Materiel
 - f.* Training

The head of MPE is the one who deals directly with the regulatory authority, so as chairman of the Maintenance Program Review Board, he or she would coordinate any recommended changes requiring regulatory approval.

The MPRB meets monthly to discuss the overall status of the maintenance reliability and to discuss all items that are on alert. The permanent members, or their designated assistants, attend every meeting; the advisory members attend those meetings where items that relate to their activities will be discussed. Items coming into alert status for the recent month are discussed first to determine if a detailed investigation by engineering is needed. Possible problems and solutions may be offered. If engineering is engaged in or has completed investigation of certain problems, these will be discussed with the MPRB members. Items that are currently in work are then discussed to track and analyze their status and to evaluate the effectiveness of the corrective action. If any ongoing corrective actions involve long-term implementation, such as modifications to the fleet that must be done at the "C" check interval, the progress and effectiveness of the corrective action should be studied to determine (if possible) whether or not the chosen action appears to be effective. If not, a new approach would be discussed and subsequently implemented by a revision to the original engineering order.

Other activities of the MPRB include the establishment of alert levels and the adjustment of these levels as necessary for effective management of problems. The rules governing the reliability program are developed with approval by the MPRB. Rules relating to the change of maintenance intervals, alert levels, and all other actions addressed by the program must be approved by the MPRB. The corrective actions and the subsequent EOs developed by the engineering department are also approved by the MPRB before they are issued.

Reliability program document

The Maintenance Review Board (MRB), derived from the FAA Advisory Circular AC 121-22B, provides guidelines for the aviation industry to use minimum scheduled interval/tasking requirements for derivative and/or newly type-certificated aircraft and their power plants for FAA approval. The AC 121-22B also refers to schedule interval requirements as the Maintenance Review Board Report (MRBR). After receiving approval from the FAA, an operator may generate or develop its own maintenance program for its particular type of fleet.

236 Oversight Functions

The air carrier may use this AC's provisions along with its own or other maintenance information to standardize, develop, implement, and update the FAA-approved minimum schedule of maintenance and/or inspection requirements for this program to become a final written report for each type certificate holder.

The MRB revision issued by the manufacturer is sent to the fleet maintenance manager (FMM) or dedicated maintenance person assigned by the air carrier. In some cases, this is the director of maintenance (DOM). The FMM/DOM interfaces with the aircraft maintenance and production department to advise them about the MRB program updates and revisions. The air carrier normally tracks each revision by fleet type to ensure the corrective action plan has been recommended to bring the maintenance production department into compliance. The MRB runs concurrent with the continuous analysis and surveillance system (CASS) and the reliability-centered maintenance (RCM) and is applied using the maintenance steering group MSG-3 system. The MSG-3 origination is associated with the Air Transport Association of America (ATA). The ATA coding system (detailed in Chap. 5) divides aircraft into distinct ATA units, and every ATA unit is analyzed for regulatory purposes to understand the results retrieved from the system and then passed on to an aviation industry steering group/committee. After the data has been reviewed by the steering committee and approved by the regulatory board for the MRB, the results are published as part of the aircraft maintenance manual.

This document also includes detailed discussion of the data collection, problem investigation, corrective action implementation, and follow-up actions. It also includes an explanation of the methods used to determine alert levels; the rules relative to changing maintenance process (HT, OC, CM), or MPD task intervals; when to initiate an investigation; definitions of MPRB activities and responsibilities; and the monthly report format. The document also includes such administrative elements as responsibility for the document, revision status, a distribution list, and approval signatures.

The reliability program document is a control document and thus contains a revision status sheet and a list of effective pages, and it has a limited distribution within the airline. It is usually a separate document but can be included as part of the TPPM.

FAA interaction

It is customary, in the United States, to invite the FAA to sit in on the MPRB meetings as a nonvoting member. (They have, in a sense, their own voting power.) Since each U.S. airline has a principal maintenance inspector (PMI) assigned and usually on site, it is convenient for the FAA to attend these meetings. Airlines outside the United States that do not have the on-site representative at each airline may not find it as easy to comply. But the invitation should be extended nevertheless. This lets the regulatory authority know that the airline is attending to its maintenance problems in an orderly and systematic manner and gives the regulatory people an opportunity to provide any assistance that may be required.