
Reliability Management Overview

Course No: B03-004

Credit: 3 PDH

Daniel Daley, P.E., CMRP



Continuing Education and Development, Inc.
22 Stonewall Court
Woodcliff Lake, NJ 07677

P: (877) 322-5800
info@cedengineering.com

Reliability Management Overview

Introduction

The field of Reliability Management or more directly Reliability Engineering has been in existence for a number of years but remains misunderstood by many individuals in the industry. You will find a number of individuals assigned to positions in which they carry the title of Reliability Engineer when in fact they perform the role of a Structural Engineer, a Rotating Equipment Engineer, an Electrical Engineer or an Instrumentation Engineer. While the functions are not totally exclusive, engineers in the more conventional roles tend to focus on the “functionality” or “integrity” of an asset rather than the reliability, availability or maintainability of an asset.

While a comprehensive treatment of the entire field of reliability would take more time than illustrated herein, it is the objective of this course to provide a high level overview of the subject. At the conclusion of this course, the student will understand the difference between the role of a true Reliability Engineer and engineers in more conventional roles. The student will also understand many of the elements that must be managed to ensure that each asset will perform in a reliable manner.

What Do You Have a Right to Expect?

If they realize it or not, when most people purchase a new asset, they have certain expectations concerning the reliability of an asset. Over the course of the life of that asset, those same individuals have continuing expectations concerning how that asset should perform.

In a commercial or industrial setting, the same is true. Senior managers of corporations have certain expectations concerning the reliability and availability of the assets they manage. As a result, they have expectations concerning how much product can be produced and how much income the organization will produce.

As with any other characteristic of a physical asset, an important question one should ask is, “what do you have a right to expect?” For instance, if you somehow remotely ordered a new car but did not specify the color and when the new car was delivered, it was bright yellow; is there a basis for a complaint? If color was not specified, one might assume that color was not important to the buyer.

While many of the details leading to reliability performance are far more subtle than the color of a new car, the steps leading to assurance that they exist are equally black-and-white. If the owner wants an asset to perform with a specific level of reliability, he must take the actions that will produce that performance. If he fails to take those steps,

the resulting performance is much like the color of the car described above. It will be the luck of the draw and will depend on factors the owner has chosen not to manage.

Each and every phase during the life of an asset contains situations that can lead to either good reliability performance or to poor reliability performance. If the owner wants to ensure good reliability performance, it is important that he takes steps during those situations in a way that will ensure the level of reliability performance he hopes to achieve. After providing some of the background needed by the student to understand details that will be later described, this course will provide brief descriptions of the steps that owners should take to ensure the desired level of reliability performance.

Definitions

There are a number of important definitions associated with the study of reliability. The definitions are critical in understanding the subtle elements that determine reliability performance and the importance of the tools used to ensure adequate performance for each of those elements.

Let's begin with the definition of the word Reliability. In this context, I will use the word "Reliability" (capital R) to identify the definition of the term most people are thinking of when they say "reliability". I will use the word "reliability" (small r) when referring to the technical definition of the specific property of asset reliability.

When most people use the term Reliability, they are actually thinking of a characteristic that involves several distinct characteristics. When discussing "good reliability" or "bad reliability" they are actually thinking about the widest spectrum of characteristics associated with the characteristic that can either cause them confidence or costs.

When using the term Reliability, most people are actually thinking of a characteristic that contains the characteristics of reliability, availability and maintainability. Each of these characteristics is distinct and each is the product of separate activities. On the other hand, each of these characteristics has a bearing on the others and therefore should be discussed together.

To define "reliability", it is a measure of the likelihood that a device will avoid failure over a specific interval of time. As a result, reliability is a statistical measure that results from understanding the actual number of failures that an entire population of a device can be expected to endure in a given interval.

It is important to understand that the actual reliability of a component is not the same as the reliability of a complete asset. The reliability of a component can vary based on how it is applied and how it is used. The reliability of a device is typically associated

with one or more specific Failure Modes that determine how and when the device will fail.

Another term using the word reliability is the term “Inherent Reliability”. The Inherent Reliability of an asset, is the likelihood that the entire asset will survive without failure over a specific period of time. The Inherent Reliability of a complete asset is based on the configuration of the asset as well as the individual reliability of the components used to construct that asset. For instance, if the configuration of an asset includes redundant components in highly critical locations, it is likely that the Inherent Reliability of the asset will be higher. Also, if components with higher individual reliability are selected over less expensive components with lower reliability, it is likely that the asset will have a higher Inherent Reliability.

The Inherent Reliability of an asset defines the upper limit or the maximum reliability performance the asset can achieve. Achieving the full Inherent Reliability requires that the asset be operated and maintained as well as possible. If the asset is operated or maintained in a sub-optimum manner, it will not be possible to achieve the full Inherent Reliability of the asset.

Another term that most people roll into their intuitive definition of Reliability is the characteristic of availability. Availability is a measure of the portion of time an asset is able to perform its intended function. The total availability of an asset is typically reduced by two factors:

1. Availability is reduced by the amount of time required to recover from unplanned failures. This portion of lost availability is dependent on both: the unreliability (or frequency of failures) and on the owner's ability to respond to the failures in a timely and effective manner.
2. Availability is also reduced by the amount of time the asset is shut down to perform planned predictive or preventive maintenance.

In simple terms:

$$\text{Availability} = \text{Total Time} - (\text{Planned Down Time} + \text{Unplanned Down Time}) / \text{Total Time}$$

Where,

$$\text{Unplanned Down Time} = \text{Sum (Unplanned Failures} \times \text{Time to Respond to each failure)}$$

Typically, major assets will require some form of major maintenance event at a number of points over the life of the asset. These major maintenance events are called overhauls, turnarounds or outages. Because these major maintenance events take so long and cost so much money, they are a major concern.

Within each major asset, there are one or more components that tend to determine how frequently the outages will need to take place. There are also one or more components that determine how long the asset will be shut down for repairs or renewal. We will call the component that determines the maximum length of time between outages “run-limiters” and the components that determine the amount of time the asset must be shut down “duration-setters”.

The “run-limiter” is typically a wearing component that becomes worn to the point that it can no longer perform its intended function. By analyzing the asset and identifying the “run-limiter” it is possible to make that component more robust or capable of enduring more wear and thus extend the period of time between outages.

The “duration-setter” is typically a component that is either buried deep in the asset or one that requires a time consuming renewal process that entails a critical path duration longer than any other component in the asset. Again, by identifying the “duration-setter” it is possible to redesign the asset in a manner that reduces the critical path duration and therefore the amount of time the asset is out of service.

The other characteristic mentioned as an element that many people include in their intuitive definition of Reliability is maintainability. Maintainability is a measure of the ability to restore the Inherent Reliability in a ratable period of time. There are two characteristics important to maintainability. The first is the ability to restore the Inherent Reliability and the second the ability to do so in a ratable period of time. A ratable period of time is a known or repeatable amount of time.

The characteristic of maintainability is easiest described by showing how one might perform a maintainability review of a new asset.

Each component of a new asset has a specific reliability based on one or more specific Failure Modes and an expected usable life. During the usable life of the asset, each component will require both proactive maintenance and reactive maintenance. The proactive maintenance consists of the predictive and preventive maintenance tasks needed to minimize unplanned failures by preventing deterioration and to restore the component to ‘good as new” conditions at the end of its useful life. The reactive maintenance consists of repairs needed to restore asset functionality and Inherent Reliability after an unplanned failure.

Once all forms of proactive and reactive maintenance that will be needed over the entire life of an asset, it is possible to review those tasks to see if they are maintainable. If the tasks include steps that are of an unsure duration or that produce uncertain results, the task is not “maintainable”. An example of a task of unsure duration is one that requires the mechanic or technician to work in an unsafe or awkward position. An example of a task that will produce uncertain results is one that concludes without a

functionality test or one that contains a step requiring an attachment with an adhesive that requires special conditions to cure.

Knowing that specific tasks that will be required over the life of an asset are unmaintainable, this gives the reliability engineer the opportunity to redesign the asset to produce a product that is maintainable thus ensuring the desired reliability and availability.


Returning to the introduction of this short course, the student should think about the assets with which he is familiar. How many of those assets have been exposed to a comprehensive maintainability review? How many have been analyzed to identify the “run-limiters” and the “duration-setters” or have been redesigned to increase the availability? Absent the maintainability review, how is it possible to ensure the asset can perform at the desired reliability and availability over its entire life? Also, look around at your current organization, who is expected to perform the maintainability review described above? While some organizations have individuals with the title “Reliability Engineer”, few have constructed roles for individuals in those positions that will ensure reliability, availability and maintainability performance at specific desired levels.

Patterns and Relationships

Like so much of engineering, the management of Reliability depends on observation of patterns of events and the relationship of those patterns of events with failures. Unlike many other engineering disciplines, the observations of patterns and relations have not been established and codified by individuals with names like Newton, Planck, Bernoulli, Ohm and others found in text books. In the business or reliability of your assets, you are the one who will need to record information and analyze it to identify the patterns and relationships leading to the failures of your assets. Even if two of the same assets were purchased on the same day by two different companies, they would be likely to have different Reliability performance. That is because no two companies use their assets in exactly the same way. As a result, the patterns of events leading up to a failure and the usable life and failure modes are likely to differ. For that reason, it is important that the Reliability Management Process and the Reliability Engineers become experts on the patterns and relations specific to their own company.

One of the well known ways of describing the relationships between a specific pattern of events and an associated failure is a diagram describing the P-F Interval of a specific Failure Mode. In this context, the term “P” refers to the earliest point that the potential for failure is known to exist. The term “F” refers to the failure event or the point at which the component in question has experienced the amount of deterioration needed to produce a failure.

The chart below is intended to describe the P-F interval starting with the point when the Total Base Number of a lubricant has become too low and the point when a related failure occurs.



The following are the definitions of the sequential elements on the chart:

- TBN – Normal (Inherent or additive based “reserve alkalinity” or ability to neutralize acidity)
- TBN – Marginal
- TBN – Low
- Source of acidity introduced
- TAN – Increases (Acidic concentration of oil)
- pH – Reaches a point where component deterioration is possible – As evidenced by some other measure
- Component Deterioration
- Deterioration to the point of added costs to repair or replace components
- Deterioration to the point that engine failure is possible

TBN is Total Base Number of the oil sample,

TAN is the Total Acid Number of the sample.

When analyzing the P-F chart in general, it is possible to say the following:

- As long as the TBN is normal, there is little or no risk of damage to the system as a result of acid contamination.
- When a marginal TBN is detected, there is a warning that something is consuming the reserve alkalinity. There is no immediate concern over acid related damage because with reserve alkalinity the lubricant cannot become acidic.
- When TBN is low, there is an immediate concern over introduction of acidity.
- The introduction of a source of acidity is typically not something under the direct control of the operator, so this event can happen at any time.
- Unless the oil sampling process is on a frequent and highly structured routine, finding a sample with high TAN can be a hit-or-miss process.
- Once the pH of the oil becomes too low, any lubricated surface can be subject to acid attack.


- Depending on how much acid exists in the system and how heavily the asset in question is currently loaded, deterioration to sensitive components can begin quite soon.
- At some point the amount of damage to components will add to the repair costs.
- Finally, depending on the amount of deterioration and the current load on the asset, a catastrophic failure can occur.

While the pattern described above is generic in nature, the timing and degree of the elements it contains is specific to each owner. For instance, one manufacturer of locomotive engines produces only four cycle engines and the other produces only two cycle engines. The likelihood of acid-gas bypassing from engine exhaust gasses are different for the two kinds of engines. As a result, the relative timing and associated risk of the P-F Interval for these two types of engines will differ. The manner in which a reliability engineer for a company using one type of engine will be different than a reliability engineer for a company exclusively using the other type of engine.

While the generic pattern and relationships for an entire industry may be very similar, ultimately, the patterns and relationships are very much company specific. A significant portion of Reliability Management and Reliability Engineering is understanding the patterns and relationships important to their industry. Even more important is applying that knowledge to the specific patterns of their own company and plant.

In the example provided above, the rate at which the total base number may deteriorate may be similar among a variety of different users. On the other hand, the timing at which the source of acidity is added may vary significantly. By way of example, there are two major manufacturers of freight locomotives. The engines in the units provided by one manufacturer are two-stroke engines and the other manufacturer produces four-stroke engines. This difference will result in some differences in the rate at which the lubricating oils are exposed to blow-by exhaust gasses; the most common source of acidity in engines. As a result, the shape of the P-F curve and the timing of the relationship between TBN deterioration and engine damage would be different in these two cases. The reliability engineer working for a company using primarily two-stroke engines would need to respond to this pattern and the associated relationship differently than the reliability engineer working for a company using primarily four-stroke engines.

The following is another, far more common example that uses oil as a basis for understanding patterns and relationships and how they may vary from situation to situation:



In this case, the patterns begin with easily controllable activities including:

- Timely oil changes
- Timely oil sampling
- Routine checking of oil temperature by sensing the temperature of the bearing housing
- Routine observation of the oil color by viewing the oil bulb on the side of the bearing housing

The pattern of events leading to ultimate failure continues with the following steps:

- Deterioration or thinning of the oil to the point that the film thickness between the shaft and the bearing is no longer adequate.
- Damage to the rotating shaft or the bearing
- Damage to the overall pump where it will no longer perform its intended function
- Catastrophic failure of the pump leading to a plant fire.

The value of understanding the pattern of events and the relationship between early signs of deterioration and events leading to failure is that the reliability engineer can describe how best to respond to them.

Timely Oil Change	How determined?
Timely Oil Sampling	How determined?
Check Bearing Housing Temp.	Who does it and when?
Check Oil Color in Oiler	Who does it and when?
Oil Film Thickness	How can this be done?
Bearing Damage	How Detected?
Shaft Damage	How Detected?
Pump Damage	How Detected?
Fire - Plant Damage	How Detected?

The first four items on the list are relatively simple to manage. The first two, it is necessary only to determine the most cost effective timing to perform the tasks. The second two items depend on deciding who in the organization has the best opportunities (available time and ready access to make observations).

The fifth item, determining the oil film thickness in a timely manner when it begins to thin is impossible. The last four items are activities that have to occur at a time when the reliability engineer has no control over. Once bearing damage has begun, the following events can happen instantaneously or they can take quite some time to occur. As a result, the P-F curve for this particular situation tells us that we need to act on the activities over which we have control. We need to change oil on time, sample it on a regular interval, monitor temperature regularly and monitor color on frequent, regular intervals.

While the two examples provided above are very specific in nature, the concept of being observant and identifying patterns leading to deterioration and their relationship with ultimate failures is one that can be applied to numerous situations in an industrial or commercial setting.

Roles in Reliability

While there are a number of roles in a typical maintenance or reliability organization, there are several that should be emphasized as to their importance in relation to performing the day-to-day activities involving processing information and investigating failures.

Diagnostician – In the process of processing the activities that occur as the result of a system failure, one role that is important to the effective and efficient use of information is the Diagnostician. While this might be a single job in a large organization, it might also be a shared job in a smaller organization. The important point is that the activity of performing a diagnosis is one that is separate and distinct.

A diagnosis is an activity that can be performed remotely based solely on the content of a properly structured Malfunction Report and the information contained in historical files and current operations data. Based on that information, a diagnostician should be able to identify the most likely Failure Mode as well as other possible Failure Modes in ranked order. The diagnostician provides instructions to the troubleshooter who is next involved in the chain of events leading to a repair.

Troubleshooter – Troubleshooting is an invasive activity that entails disassembly of the failed system to identify the failed component and the condition of that component. A troubleshooter should be provided with instructions concerning where to start by the diagnostician. Starting disassembly without knowing the location of the most likely Failure Mode is problematic, because it leads to wasted time and occasionally, introduction of new defects into the system that did not exist before the troubleshooting began.

The troubleshooter identifies the failed component and provides a detailed description of all the actions needed to perform a complete and thorough repair.

Failure Analyst – Once the repairs are underway or complete, another step accomplished by an organization interested in learning and preventing future failures in Failure Analysis. Failure Analysis is the step that identifies the Failure Mechanism. While the Failure Mode is the result of deterioration that has been occurring for some time up to the failure, the Failure Mechanism is nature's tool for forcing the deterioration to occur.

In mechanical components, there are four forms of Failure Mechanisms:

- Corrosion
- Erosion
- Fatigue
- Overload

As described above, there mechanisms causing deterioration exist in nature and if the techniques used to prevent them are not maintained, the deterioration will be allowed to proceed unabated until the failure Mode is present and a failure can occur.

For instance, a protective coating may prevent uniform corrosion. Absent, the presence of the coating, uniform corrosion can occur resulting in metal loss, thinning and ultimate failure. Another example is the rubber seal that keeps moisture out of an electric cabinet. Inside the cabinet are a variety of dissimilar metals. If the seal is not maintained, water can intrude setting up a battery between more noble and less noble metals. Ultimately this can either result in metal loss leading to failure or it can produce

sufficient corrosion products that will find their way into the contacts and cause further deterioration and failure.

In order to understand the Failure Mechanism that led up to the Failure Mode, it is important for a Failure Analyst who is familiar with the various forms of Physics of Failure to investigate all significant failures. One point to keep in mind is that if some form of Failure Mechanism is actively working in one part of your asset, it is likely working in other parts. Identifying and eliminating that Failure Mechanism wherever it exists can prevent a number of failures; not just the one that was investigated.

Cause Analyst – While most failures have a variety of causes, there are at least three levels of cause:

- Physical Cause
- Human Cause
- Systemic Cause

At the lowest level, there is always one or more Physical Causes that unleash nature's Failure Mechanism so deterioration can begin. In the cases described above, the absence of the protective coating or the absence of the rubber seal is the Physical Cause that allowed the direct contact of water with the metal surfaces. In each case, that resulted in the Failure Mechanism of Corrosion starting the process of deterioration.

At a level above the Physical Cause is the Human Cause. The Human Cause is a specific person who either acted or failed to act in a manner that resulted in the Physical Cause. In the case of the rubber seal on the electrical enclosure, several individuals may be the physical cause:

- If the maintenance work order called for maintaining the rubber seal at some point and the work was not done, the crafts person who was assigned to perform the work order may be the Human Cause.
- If the electrical engineer assigned to create the maintenance work order failed to identify the need to regularly inspect and maintain the rubber seal, he might be the Human Cause.
- If the manager over the area where the electrical enclosure is located decided to save some money by removing the seal maintenance task from the work order, he would be the Human Cause.

In any case, it is imperative that the specific person who is the Human Cause be identified. Speaking to that person is the only way the Systemic cause will be identified.

The Systemic Cause is best described as a “trap” that exists in the organization, the procedures, the accepted practices or the behaviors of the overall organization that allows the Human Cause to act in a manner that produces the Physical Cause.

In the examples described above:

- If the crafts person can pick and choose what portions of the work order he wants to complete, the culture has created a trap that has led to this failure.
- If the electrical engineer does not have the time or has not been provided the guidance to get out in the field and identify the need for replacing the seal, there is another systemic weakness.
- If cost cutting and budgetary restraints have caused the manager to remove needed work from the work order, there is a different systemic cause.

In any case, the Systemic Cause that is identified can affect a much broader array of issues than just the one being investigated. The person performing Cause Analysis can either be the same person or can be a different person than the Failure Analyst. The two roles typically require different skill sets and are accomplished in different settings, so a single individual assigned to both roles typically does one better than the other.


Governing Patterns in the Life of a Single Failure

The descriptions provided above tended to focus on specific details concerning either the patterns and relations leading to failures or the roles and skills involved in a reliability organization.

This section will provide the first of two general patterns that are important to understand when creating systems that will properly address failures, gather information on failures and take the steps needed to ensure a long and reliable life for an asset.

The first of the general patterns to be described is one that identifies the steps that occur in either an overt or passive manner during the life of a single failure event. If handled in the proper manner, these steps can lead to effective and efficient handling of the single event. If handled well, these steps can also lead to a permanent solution to the problems that caused the failure to occur. Handled poorly, the defect causing the immediate event may not be removed. This can lead to a repeat of the failure and a general deterioration of the inherent reliability of the system.

The title used herein for this pattern is the Path to Failure and Corrective Action because it identified all the discrete steps leading to a failure and resolving it. While many organizations may not openly acknowledge the presence of each step in the process as it is described, the steps exist and are being handled in a passive manner without specific focus or intentional management.



As described above, any failure event typically begins with a Systemic Cause. The Systemic Cause creates a “trap” into which an individual can fall. The next step is the Human Cause in which a specific individual either does something or fails to do something that produces a Physical Cause. The third step is the Physical Cause. This is the step in which the form of prevention used to restrain nature's Failure Mechanisms are ignored.

The next step is the Failure Mechanism. The Failure Mechanism (like corrosion, erosion, fatigue or overload) is a natural process that causes on-going deterioration. At some point in time, the amount of deterioration has reached the point that a defect is present. The defect is not the same as the failure. In the case of corrosion, the defect might be the point at which deterioration has removed enough metal that the component is no longer able to handle the load at its maximum condition. Once the defect is present, the failure clock is ticking and the failure is waiting only on a situation when the loading is sufficient to cause a break at the point of maximum deterioration. In the case of a rusted pipe, the leak may not occur until the system pressure is raised to higher than normal pressure.

Once the failure occurs, the next step is the Malfunction Report. In some cases, the structure used for issuing a Malfunction Report ensures that clear and accurate information is provided. In other cases, the information provided is useless and provides little help in driving toward an accurate diagnosis and an ultimate solution.

The next three steps tend to work together in the effort to identify the Failure Mode and the ultimate resolution. The Diagnosis is the hands-off, remote activity of using available information to identify likely Failure Modes. This information can be used to perform triage or prioritizing the handling of current issues. It can also be used for advanced preparation of tools and materials. The next step is “funneling” or analysis of all possible Failure Modes to determine which should be approached and in what order. The most likely Failure Mode is typically first to be approached by the troubleshooter. On the other hand, there might be a relative low likelihood Failure Mode that requires little time or effort to “reset”. That item might be attempted first on the off-chance it produces immediate resolution. Recycling a computer is an example of such a fix. Troubleshooting is the third of these activities. Since it is hands-on and invasive, it takes the most time and it can also introduce new defects into the system. Highly directed troubleshooting typically leads to the most effective and efficient repairs.

The next two items also fit together in terms of identifying the Failure Mode that caused the failure. The troubleshooter should identify the defective component that caused the failure. He should also describe the condition of that component that lead to the defect. Frequently replaced components that restore functionality are not defective. Simply getting inside and shaking things up restores functionality by restoring poor contacts. In this case, the replaced component is not defective. In this case, the true defect has not been removed and the inherent reliability of the system has not been restored. Thus it is critical that the troubleshooter produce a truly defective component.

It is also important that the troubleshooter describe the condition of the defective component. Without such a description, it is impossible to identify the Failure Mechanism. Without knowing the Failure Mechanism, it is impossible to eliminate the natural process forcing deterioration.

The next two steps in the process are Failure Analysis and Cause Analysis. As described earlier, Failure Analysis is an analysis of the physics of failure to determine the source of the deterioration leading to failure. Cause Analysis is the humanistic and organizational part of the Corrective Action process.

The chart provided above includes two other elements. The first is identification of Failure Mechanisms that are in the process of producing deterioration before failure. It is important to note that the Failure Mechanism is hard at work for a long period before the component is deteriorated to the point that failure is possible. If the Failure

Mechanism is identified and remedied before the failure, it is possible to prevent the failure from occurring.

The second element highlights the opportunity to identify the defect after it has formed but before the failure has occurred. As mentioned earlier, the formation of the defect and the failure are not always simultaneous. If an alert person can find a defect that exists before a failure, it is also possible to prevent the failure from occurring.

When each of these steps are recognized and handled correctly, it is possible to:

- Handle each individual incident in an effective and efficient manner
- Gather the information needed to handle future incidents in an effective manner
- Gather information needed to prevent this incident in the future
- Gather information needed to prevent similar incidents caused by the same Failure Mechanism or Failure Mode

Investigation of the typical patterns of response to individual failures and how those patterns ultimately result in either permanent solutions or reduced inherent reliability will be useful in identifying the actions that must be taken to improve reliability performance.

Governing Patterns in the Overall Life of an Asset

The overall lifecycle of an asset can be broken into a series of processes that ultimately determine the reliability, availability and longevity of the asset. Those processes include the following:

- Design and build – The inherent reliability of an asset is based on the configuration of the asset as well as the individual reliability of the components that were selected. Many conventional design processes address only the functionality and the structural robustness of an asset. While necessary, these forms of analysis do little to ensure the reliability, availability or maintainability of an asset. Also, they do not determine the usable life of the asset. In order to address those characteristics, it is necessary to perform Design for Reliability concurrently with conventional design activities.
- Operate – Assuming that an asset has been designed in a manner that the Inherent Reliability is capable of providing the performance desired by the owner, that is only the beginning. In order to provide the required performance, the asset must be operated and maintained in a manner that harvests all the

Inherent Reliability. There are two things that an operator of an asset can do to support good reliability.

First, the operator can operate the asset in a manner that he will “Do No Harm”. If the operator understands the Failure Modes and the Failure Mechanisms that may deteriorate the equipment he operates and he understands the choices he can make that will either prevent or cause deterioration, he can choose the path that will reduce the amount of harm that results from poor operation.

Second, the operator can “Do Some Good” as part of his operating practices. In earlier discussions of P-F intervals, the opportunity for someone to regularly observe oil temperature and oil color were mentioned. Operators are typically the resource with the greatest and most frequent opportunity to make these observations. If the operator takes action when the bearing housing of equipment items are too hot or when the oil is just beginning to discolor, it is possible for the operator to “Do Some Good”.

While the examples provided are fairly simple, there are myriad examples for operators to modify the way they interface with equipment to positively affect the equipment reliability so deterioration is avoided and the full Inherent Reliability is achieved.

- Inspect – Throughout the life of an asset, there are situations in which individuals with specialized expertise perform inspections of the equipment. Frequently the individuals are trained to recognize corrosion, vibration, electrical system deterioration and a variety of other forms of specialized patterns that provide clues pertaining to incipient or on-going deterioration. These individuals are made more effective if they are made aware of Failure Modes and Failure Mechanisms known to be present in the actual systems being inspected. Rather than looking for everything, they can focus their attentions on problems that have happened in the past and are likely to occur again in the future.
- Maintain – One of the on-going routines during the lifecycle of an asset is the need to perform maintenance. There are two forms of maintenance: proactive and reactive. Proactive Maintenance can either be Predictive or Preventive. Predictive Maintenance is typically non-invasive and uses special tools and techniques to identify signs of incipient failure or on-going deterioration. Activities performed as part of an inspection is typically intended to be predictive in nature. Preventive Maintenance is another form of Proactive Maintenance. Preventive Maintenance is tangible activities performed to exchange or renew a component based on knowledge that time has come for the replacement to occur.

Reactive maintenance is the form of maintenance that is accomplished in response to a failure. The Path to Failure and Corrective Action described above provides a comprehensive description of the steps included in reactive maintenance.

In either case (proactive or reactive maintenance) the objective is not to simply restore or ensure the functionality of an asset. The objective is to maintain or restore the inherent reliability of the asset.

- Overhaul, Turnaround or Outage – On some regular but non-routine basis major assets need to be maintained using a major effort that is called an overhaul, outage or turnaround depending on the industry. These events are highly costly and they have the single greatest impact on the availability of the asset involved. Most systems have either one or a small number of components that reach the end of their useful life and, as a result, set the timing when the overhaul, outage or turnaround must occur. For simplicity we can call these “run-limiters”. There is also one or a small number of components that require the longest critical path of activities during the overhaul, outage or turnaround. Again for simplicity, we can call these items the “duration-setters”.

Apart from reducing the number of failures and thus increasing the amount of unplanned outage time, the best way to improve the availability of an asset is to:

1. Make the “run-limiters” more robust so the asset can run longer between outages
 2. Make the “duration-setters” more maintainable so the duration of outages is shorter
- Modification – Major assets are frequently modified over their lifecycle. They can be modified to alter their performance or to increase their capacity. It is not uncommon for modifications to be accomplished without adequate attention paid to the Design for Reliability. In those cases, the resulting reliability of the modified asset is less than the asset before modification.
 - Renewal – As with modification, great many assets go through a renewal process to “breath new life” into aged assets. As with modifications, it is not uncommon for inadequate attention being paid to Design for Reliability during the renewal process. In this case, the asset is once again prepared for a long but unreliable life.

An important point to keep in mind is that it is critical to remain vigilant over the entire lifecycle of an asset to ensure good reliability. You cannot be vigilant 90% of the time

then drop your guard. The damage done during even a short period of inattention can result in a loss of critical characteristics that it will be difficult or impossible to replace.

The importance in understanding the pattern of events and processes that occur over the entire lifecycle of an asset comes in being able to determine if reliability is being properly assessed at those times. For instance, there are organizations that have reliability analysis well integrated with the processes where the engineering discipline is involved. Since initial design, modifications and renewal typically involve engineers in the activity; those three activities would be the ones most likely to benefit from that involvement. In some organizations, the opposite can also be true. Activities handled by the plant resources (operation, maintenance and inspection) may benefit from the involvement of reliability engineers in those activities while project managers who are focused solely on cost and schedule may refuse to address additional requirements that add cost or take more time.

For an asset to be truly reliable, it is important that elements and activities leading to good reliability are addressed on all occasions.

Tools for Each Phase in the Life of an Asset

Each point in the life of the asset has activities that can help improve the reliability, availability and maintainability of an asset. On the other hand if these characteristics are ignored at any of those points, the characteristics needed to produce reliable performance can be lost. The following describes the tools typically associated with each phase in the lifecycle of an asset:

- Design and build – During the design and build process, it is important to use a comprehensive Design for Reliability (DFR) process. The DFR activities need to be accomplished concurrently with conventional design activities. If DFR activities tend to lag conventional design, components will be chosen and purchased and it will be difficult to make changes.

One of the key elements of the DFR process is a tool called Reliability Block Diagram (RBD) analysis. During RBD analysis a model of the ultimate system design is created. Each major component that is subject to failure is represented by a single box and each box contains the factors that represent the statistical reliability of the element it represents. After the model is assembled, it is then possible to calculate the system reliability using either repeated simulations or manual calculations. In either case, the results provide an estimate of the expected reliability of the system. If the calculated or simulated reliability is less than required it is possible to add redundancy to specific locations or to increase the reliability by selecting more reliable components. In either case, re-running the calculations or simulations will show if the modifications are adequate or if additional changes need to be made to achieve the target reliability.

Another step in the **DFR** process is to calculate the availability of the system. This can be done in two ways. First, if **RBD** software has been used to calculate the expected reliability, it is typically possible to add information describing the expected response time to failures and describing the anticipated amount of planned outage time to account for Predictive Maintenance, Preventive Maintenance and Outages for Overhaul or Turnaround.

It is also possible to perform a manual estimate of the availability. This is done by identifying the longest cycle between non-repetitive events then identifying the total down time during that complete interval.

As an example, assume that a plant requires a minor outage for boiler inspection each year. The outage period for these outages is one week. Also assume that the plant requires a two week outage for catalyst change every other year. In alternating years, the boiler inspection can be done during the catalyst renewal. Finally assume that every ten years a six week long turnaround is required.

The cumulative downtime is as follows:

Year 1 – 1 week

Year 2 – 2 weeks

Year 3 – 1 week

Year 4 – 2 weeks

Year 5 – 1 week

Year 6 – 2 weeks

Year 7 – 1 week

Year 8 – 2 weeks

Year 9 – 1 week

Year 10 – 6 weeks

Total Down Time – 19 weeks

Total time in cycle – 520 weeks

Planned Availability = $(520 - 19) / 520 = 96.53\%$

For the sake of completeness, let's assume that the plant also experienced one unplanned failure each year and it required one week to recover from each unplanned outage. In that case the Unplanned reliability is as follows:

$$\text{Unplanned Reliability} = (520 - 10) / 520 = 98.07\%$$

The combined availability would be:

$$\text{Availability} = (520 - 19 - 10) / 520 = 94.42\%$$

If the capacity demand for the product being provided by the plant is more than 95% of the design capacity of the plant, the availability would be inadequate. It would be necessary to take steps that would reduce either the planned outage time or the unplanned outage time.

Another tool useful in performing the DFR analysis is Reliability Centered Management (**RCM**) analysis. RCM analysis is useful in identifying all the predictive and preventive maintenance tasks that will be conducted over the life of the asset. It is also useful in identifying all the elements that will be allowed to “run to failure” so the reactive repair tasks can be identified. Once the complete list of Predictive, Preventive and Reactive repair tasks are identified, it is possible to describe the steps needed to complete the tasks. By performing a mental or physical “walk through of all tasks it will be possible to determine if the asset is maintainable. If tasks require an unsure amount of time to complete or produce uncertain results, the tasks and the asset are not maintainable.

- Operate – There are two popular tools for improving the operator interface with the equipment he operates. One is called Total Productive Maintenance, the other is Operator Driven Reliability. In both cases, the objective is to use added structure and discipline in the relationship between the operator and the equipment to facilitate the objectives of doing no harm to the equipment and performing some activities that will do some good.
- Inspect – The inspection process used to identify and monitor on-going deterioration can be significantly enhanced using a well structured and disciplined **Failure Mapping** process as described in the Path to Failure and Corrective Action discussion earlier. Close tracking of Failure Modes and identification of associated Failure Mechanisms will identify the ongoing deterioration that should be the focus of inspection efforts.
- Maintain – The objective of both proactive and reactive maintenance tasks is to maintain and restore the inherent reliability of the asset. RCM is an excellent tool for identifying the tasks that will be completed over the lifecycle of an asset. Once tasks are identified and managed by the Computer maintenance

Management System (**CMMS**), it is important to ensure that all the tasks are being done in a manner that restores the Inherent Reliability. The following are examples of situations that will fail to restore the Inherent Reliability:

- Fail to maintain redundancy as included in the initial design
- Do not use replacement parts with the same robustness as original parts
- Use inappropriate procedures
- Ignore quality control and quality assurance steps at the completion of tasks
- Take short cuts
- Ignore critical tolerances, fits and clearances in assembly of equipment
- Overhaul, Turnaround or Outage – Overhauls, Turnarounds and Outages contain many of the same elements as simple maintenance above so it is useful to review that section. In addition, these major events typically are intended to provide reliable operation of an asset for much longer time than typical maintenance. As a result, it is critical that run-limiters receive special attention and that they be provided with sufficient wear allowance to ensure they will survive for the entire intended run-length.
- Modification and Renewal – Both Modification and renewal contain many of the same elements of the initial project design. If new choices are being made concerning changes to the configuration of the asset or choices of the reliability of replaced components, RBD will be helpful in making the decisions.

Two additional tools that will assist in making sound decisions during modifications or renewal activities are Lifecycle Costing (LCC) and Total Cost of Ownership (TCO). LCC takes into consideration all costs that will occur over the entire lifecycle of the asset.