Introduction to
R & M Management

Fred Schenkelberg

Carl S. Carlson

Fred Schenkelberg
Reliability Engineering and Management Consultant
FMS Reliability
www.fmsreliability.com
fms@fmsreliability.com
@fmsreliability

Carl S. Carlson
Reliability and FMEA Consultant
ReliaSoft Consultant
www.effectivefmeas.com
carl.carlson@effectivefmeas.com

SUMMARY & PURPOSE

The purpose of this tutorial is to introduce an outline to guide the management of an effective reliability or maintainability program. Reliability, maintainability, availability, or the 'ilities' are common in our language with reference to products, services, equipment, and people. Joe is regularly available for the meeting; We can count on (depend or rely) Sara to finish the report on time; My car starts every morning without fail; and many more. What is meant with these concepts and specifically how do we manage achieving and sustaining business objectives related to these 'ility' concepts? The purpose of this short paper is to provide an introduction to key concepts and approaches commonly used for reliability and maintainability management.

With some common sense, an appreciation of the goals, understanding of expected and past failures, and the proper application of reliability engineering tools, you can manage to improve profitability, increase throughput, or enhance a brand image. With a sound design, robust supply chain, consistent manufacturing, and adequate maintenance nearly any product or complex system can meet or exceed their reliability or maintainability goals.

Fred Schenkelberg

Fred Schenkelberg is a reliability engineering and management consultant with Ops A La Carte, with areas of focus including reliability engineering management training and accelerated life testing. Previously, he co-founded and built the HP corporate reliability program, including consulting on a broad range of HP products. He is a lecturer with the University of Maryland teaching a graduate level course on reliability engineering management. He earned a Master of Science degree in statistics at Stanford University in 1996. He earned his bachelors degree in Physics at the United States Military Academy in 1983. Fred is an active volunteer as the Executive Producer of the American Society of Quality Reliability Division webinar program, IEEE reliability standards development teams and previously a voting member of the IEC TAG 56 - Durability. He was a Senior Member of ASQ and IEEE. He is an ASQ Certified Quality and Reliability Engineer.

Carl S. Carlson

Carl S. Carlson is a consultant and instructor in the areas of FMEA, reliability program planning and other reliability engineering disciplines. He has 30 years of experience in reliability testing, engineering, and management positions, and is currently supporting clients of ReliaSoft Corporation with reliability and FMEA training and consulting. Previous to ReliaSoft, he worked at General Motors, most recently senior manager for the Advanced Reliability Group. His responsibilities included FMEAs for North American operations, developing and implementing advanced reliability methods, and managing teams of reliability engineers. Previous to General Motors, he worked as a Research and Development Engineer for Litton Systems, Inertial Navigation Division.

Mr. Carlson co-chaired the cross-industry team that developed the commercial FMEA standard (SAE J1739, 2002 version), participated in the development of SAE JA 1000/1 *Reliability Program Standard Implementation Guide*, served for five years as Vice Chair for the SAE's G-11 Reliability Division, and was a four-year member of the Reliability and Maintainability Symposium (RAMS) Advisory Board. He holds a B.S. in Mechanical Engineering from the University of Michigan and completed the two-course Reliability Engineering sequence from the University of Maryland's Masters in Reliability Engineering program. In 2007, he received the Alan O. Plait Award for Tutorial Excellence. He is a Senior Member of ASQ and a Certified Reliability Engineer. His book, *Effective FMEAs*, was published in 2012 by John Wiley & Sons.

Table of Contents

1. INTRODUCTION

Reliability, maintainability, availability, or the ‘ilities’ are common in our language with reference to products, services, equipment, and people. Joe is regularly available for the meeting; we can count on (depend or rely) Sara to finish the report on time; my car starts every morning without fail; and many more. What is meant with these concepts and specifically how do we manage achieving and sustaining business objectives related to these ‘ility’ concepts? The purpose of this short paper is to provide an introduction to key concepts and approaches commonly used for reliability and maintainability management.

With some common sense, an appreciation of the goals, understanding of expected and past failures, and the proper application of reliability engineering tools, you can manage to improve profitability, increase throughput, or enhance a brand image. With a sound design, robust supply chain, consistent manufacturing, and adequate maintenance nearly any product or complex system can meet or exceed their reliability or maintainability goals.

Reliability is a quality aspect of a product. Or, as some like say, reliability is quality over time. Either way, the basic definition we will use here is reliability is the probability of a product successfully functioning as expected for a specific duration of time within a specified environment. For example, a TV remote control has a 98% probability of successfully controlling the associated television (change volume, channels, etc.) for two years in a North American home environment.

There are four elements to the reliability definition: 1) Function, 2) Probability of success, 3) Duration, and, 4) Environment.

Maintainability is related to reliability, as when a product or system fails, there may be a process to restore the product or system to operating condition. Maintainability is a characteristic of design, assembly, and installation that is the probability of restoration to normal operating state of failed equipment, machine or system within a specific timeframe, while using the specified repair techniques and procedures. We often consider two other ‘ilities’ with maintainability: 1) serviceability or the ease of performing inspections, diagnostics and adjustments; and, 2) reparability or the ease of restoring functionality after a failure.

Closely related to reliability and maintainability is availability. Availability is a characteristic of a system (piece of equipment or product) to function as expected on demand. One way to measure availability is the percentage of time the system is functioning per year. An example may be the cable TV service to a home is 95% available, meaning that for the 100 hours of desired TV entertainment per year the system functioned for only 95 hours, and was not functioning (under maintenance, power outage, etc.) for 5 hours.

The fundamental idea from a customer’s point of view is the product works as expected. The car starts, the bottling machine fills the bottles accurately and quickly, the printer just works. When asked, a customer does not want any failures, especially with the specific product that they purchase. They do want to enjoy the benefit provided by the functioning product. Unfortunately the variability of materials, assembly techniques, environments, faulty software and human errors do lead to failures. Failure happens.

In one sense, reliability and maintainability management is the management of failure. The specific approaches and tools available to the R&M manager permit the optimization of the problem to finding a cost effective solution to the design, assembly and use of a product. Reliability and maintainability engineering pulls resources and skills from across many fields including design, materials, finance, manufacturing, environmental, and statistics. R&M engineers are often asked in one form or another only two questions: 1) What will fail? and 2) When will it fail? For each specific situation (i.e. satellite or game controller) the R&M engineer assesses risk, balances the probability and consequence of failure with value, and negotiates with development, manufacturing, suppliers and customers to deliver a reliable and maintainable solution. It is often an exciting, rewarding and challenging role.

Acronyms and Notation

COGS Cost of Goods Sold
FMEA Failure Modes and Effects Analysis
FRACAS Failure Reporting, Analysis, and Corrective Action System
HALT Highly Accelerated Life Test
R & M Reliability and Maintainability

2. SPECIFICATIONS

“Faster, Better, Cheaper” is one way to state the evolving changes in many products, computers in particular. More features, more value, smaller, lighter, and of course less expensive. Add ‘lasts longer’ to this set of requirements to round out four key drivers for any product. Functions or performance is the list of what specifically the product does. This list may be quite long and detailed, and may include everything from brand logo placement to color to the power button location and size. The set of product functions is part of the reliability definition and defines the operating state and conversely what a product failure may include. While not required, a set of functions (product features) is often detailed at the start of a product development program. During product development, the design is regularly evaluated or tested and compared to the desired set of functions.

Cost may or may not be the most important consideration over the product lifecycle, yet it is often known and tracked during product development and for maintained products during use. Cost includes COGS and it may include the cost of service and repairs. It often does not directly include the cost of failure to the customer, yet that cost may be known. For example, when a deep-sea oil exploration rig is pulling a drill string due to a part failure, it may cost close to $1 million per day. Many products during design have a cost target, and it is monitored.

Time to market (or profit or volume or similar) is another common requirement placed on the product development team. This is especially true for products with a short season for sales, such as for the holiday market. Setting milestones and deadlines is a management tool to help get the product to market in a timely and coordinated manner. Like functions and costs, shipping a product on time is routinely measured.

In any program, the priority of one over the other two may make sense and may also be clearly stated as a guide to decision making. Of course there are many other considerations during product development, one of which is product reliability. In this tutorial we will focus on reliability. Similar goals can be expressed for availability and maintainability when that is appropriate for that product or system.

Reliability is the product functions as expected within a stated environment and use profile with probability of success (not failing) over a stated duration. Clearly stating the complete reliability goal is not difficult to do at the beginning of a design program. And, once stated provides a common guide for the development decision making along with reliability test planning, vendor and supply chain requirements, and warranty accrual. The goal certainly may change over the development process, as may product features, cost targets or time to market deadlines. The reliability goal is like any other product specification; it just deals with the performance over time after placed into service.

State the reliability goal such that it includes all four elements of the reliability definition. For example, a home wireless router provides 802.1n connectivity with features specified in product requirements document HWR003, in a North American home or apartment environment, with a 96% probability of still operating after 5 years of use.

3. RELIABILITY APPORTIONMENT

An extension of reliability goal setting is to break down the goal to the elements of the product. Provide a meaningful reliability objective for each of the components or subsystems. In a series system (reliability-wise) the probability of failure for each element has to be lower than for the overall system. The opposite is true for elements in parallel (reliability-wise). For complex systems the apportionment math may become more complex, yet the concept still applies.

Providing a clear and concise reliability objective to each of your design teams and suppliers provide a means to make reliability related decisions local to the element under consideration. This may influence design margins, material selection, and validation techniques.

Keep in mind that all four elements are part of the apportioned reliability goal. Often the environment and use profile will be different for different elements of a product. While the power supply may operate full time, the hard drive may often be idle and partially powered down. The location within the product may alter the temperature the elements experience. Localize the apportioned goal or at least provide sufficient information to fully articulate and act upon an apportioned reliability goal.

The process used to create the apportionment maybe as simple as an equal allocation to each element to weighted on expected or known reliability performance (predictions, models, historical, etc.). We rarely have enough information to provide perfect apportionment from the start. It will be a work in progress as the design matures, as information becomes available, and as the design is evaluated.

Setting a breakdown of the overall goal starts the discussion and thought process how every element contributes to the overall of performance of the product.

4. FEEDBACK MECHANISMS

A goal on its own is nice and generally meaningless unless compared to performance. When shooting an arrow at a target, we naturally look for the distance between the intended target and the location of the arrow. That difference provides information to the archer on adjustments to the aim of the next arrow. For product development, setting a reliability goal or any specification requires the measurement of the performance compared to the desired performance. The difference may require changing the design or adjusting the goal.

Recall the two basic questions of reliability engineering: What will fail? And, When will it fail? These form two types of feedback often used to assess the readiness of a design to meet its objectives. The two approaches are used as is appropriate for the current situation. A new technology without any field history may require an emphasis on discovering what will fail. Then shift focus to determining how long before it fails under expected use conditions.

In another situation a product and its technology, materials, and use conditions may be well known, along with the types of failures that limit the life of the product. In this case the focus may be on design changes as they impact prolonging the life of the product, with less emphasis on what will fail (as it is already known). Of course, in many situations there is a call for both approaches.

4.1 Discovery of reliability risks

What will fail is a core question facing nearly any product development or maintenance team. Henry Petroski postulates that designers create designs that avoid failure. [1] This issue might be that the design team has to know what will fail. If that is not known, than it is difficult to avoid product failures.

Thus the team’s current situation related to understanding the expected failure mechanisms plays a role in the steps to determine the expected failure mechanisms. In the situation with known failure mechanisms and the minor design changes, there is little need to ‘discover’ failure mechanisms. The focus may shift to those areas related to the changes and validation of existing failure mechanisms. Another situation may include many uncertainties related to failure mechanisms. A design change to eliminate a specific mechanism may reveal another, previously hidden, mechanism. A new material may involve exploration of how the material will react over time to the shipping and operating environment.

Discovery can use a range of tools available to reliability professionals. This may include literature searches, FMEA, and discussions with suppliers or knowledgeable researchers. The discovery may include a wide range of testing including material characterization, step stress to failure testing, and HALT.

The intent is to find the weaknesses within a design and take steps to minimize failures. For example the team may discover the material color fades quickly in sunlight, and adding a stabilizing agent may insure color fastness. A HALT may expose a faulty layout and require a redesign of the printed circuit board. Understanding and characterizing the failure mechanisms that all designs contain permit design decisions to avoid surprises in later product testing or during use.

FMEA is a tool to pool the ideas and knowledge of a team to explore the weaknesses of a product. To some this may seem like a design review and to some extent it is just that. To some it is an exploration of each designer's knowledge of the boundary to failure. Depending on the team and amount of knowledge already known, FMEA may or may not be a fruitful tool to discover product failures. It nearly always has the benefit of effectively communicating the most serious and likely issues across the team.

HALT is a discovery tool that applies sufficient stress or multiple stresses to a product to cause failure. Starting at nominal stress levels, the HALT approach then steps up increasing amounts of stress until the product no longer functions as expected. Careful failure analysis may reveal design weaknesses, poor material choices or unexpected behavior. The idea is that the failures provide knowledge on areas for improvement. A product that has the detected weaknesses resolved is more robust thus able to with withstand normal stresses and the occasional abnormal stress load without failure.

FMEA and HALT provide information about the product design and materials that to some extent rely on previous knowledge about the expected failure mechanisms. Within the FMEA team the knowledge is shared or a new question may be explored (possibly new information revealed). And HALT applies stresses that are expected to cause failure. In each case, a new product design or material may have an unknown response to an unexplored stress. Both tools serve a purpose and have proven very useful in the failure discovery process, yet acquiring more information about possible failure mechanisms may enhance both tools and the product.

Most materials and components prior to being available for use in products undergo development and characterization. Scientific literature is full of studies of metals, polymers, chemicals, ceramics and more that explore electrical, mechanical, aesthetic and more properties. The entire process is often studied from raw material to final product and may include life studies. These studies often focus on very specific failure mechanisms that limit the life of the material, assembly or component.

Modern products may have hundreds of materials and thousands of components, yet each has some history of exploration and characterization of failure mechanisms. As a minimum for new materials or components do the research to understand the known failure mechanisms and how they will behave within your design and environment. Published literature in scientific and engineering journals is a good place to start. Then engage the researchers in a discussion about what they know and how the material may behave in your design. Many component and material suppliers have intimate knowledge of the component or material weaknesses and are willing to share that with their customers. For locally invented or constructed materials or components, embark on a characterization study to fully understand the failure mechanisms.

Knowledge provides a means to understand the limiting boundaries around any design that transition the product into failure. Understanding those boundaries in your product's circumstance permits improvements to occur in areas that would otherwise lead to premature failure. Discover the failure mechanisms and how they manifest themselves in your design or system. Then you have found the answer to what will fail.

4.2 Determine duration

The second question facing a product development or maintenance team is related to how long before failure occurs. The reliability engineer, armed with knowledge around the expected failure mechanisms is in a good position to answer this question. Knowing when a product is expected to fail provides feedback to the team for comparison with the goals. It also provides a means to plan for preventative maintenance, plus contributes to spares stocking levels.

Oliver Wendell Holmes wrote a poem titled "One Hoss Shay" [2] in which a parson crafted a shay where every part was a strong as every other part. After 100 years and a day every part failed at the same time, nothing before any other. If we could create a cell phone that would last exactly 5 years, and every part failed, not one before another, we could call that perfect reliability. Not a single element of the product had any remaining usefulness. Nothing wasted.

Unfortunately, perfect reliability is difficult to achieve, as there are so many variables and unknowns related to when and how a failure occurs. In the poem there is only one shay so we'll never know if an entire fleet would also survive exactly 100 years and day.

In practice, even with literally hundreds or thousands of ways a product can fail, there are generally only a few that will dominate the initial product failure. Understanding the time until failure for these few failure mechanisms is possible for any design or maintenance team. There are three broad sources of product failure:

1. Supply chain and manufacturing variation
2. Overstress conditions during transportation or use
3. Wear out of one or more components

4.2.1 Supply Chain and Manufacturing

Even raw material suppliers use equipment such as shovels and trucks, which they procured. The ability to create a product is often reliant on the supply chain being able to provide consistent materials and components. If the material property that is important to the functioning of your product varies unacceptably, your product is more likely to fail. If the manufacturing process varies unacceptably and produces inferior product, those too are more likely to fail.

Being clear with specifications, especially concerning reliability, helps your supply chain and manufacturers create materials, components and products that will meet your reliability requirements. Besides reliability apportionment mentioned above, you may need to characterize the material properties that directly impact product reliability. For manufacturers, understanding the elements most at risk to moisture, electrostatic discharge, or corrosion and related causes of premature product failure, will assist them to create reliable products without latent defects.

Specifications, critical to reliability flags, process control, and monitoring are all tools available to the reliability professional to minimize product failures due to supply chain and manufacturing issues. A good practice when it is possible is to move the assessment and monitoring of reliability as far up stream in the supply chain as possible. Manufacturers commonly practice this, as they know building products with faulty components reduces yield and increases the cost of products. After discovering the salient failure mechanisms identify where in the process the source of the weakness may occur and control the process at that point.

The wrong material or poor assembly of a product tends to lead to early life failures. When the supply chain and manufacturing processes are working properly the unwanted variations will be identified and eliminated before a product goes to market. Those defects that make it to market may have no effect on product life or shorten product life. Predicting the impact will take understanding the nature of the variation and how that will interact with use conditions.

While difficult to predict, as the nature of the failure mechanism may be unknown, it may require study when the consequence of failure is high and the possibility of unwanted variation is high. One technique is to create products with a range of material or manufacturing variation. Then evaluate the impact on product life. This may lead to an improved product or understanding of the need to carefully control the incoming material and assembly processes. Normally, we do not attempt to predict how long products with unknown supply chain or manufacturing errors may last. The proper focus most of the time is on supply chain and manufacturing consistency and control.

4.2.2 Overstress

Mechanical engineers learn about stress versus strength when sizing a beam to carry a load. Both the beam and the load will vary from the particulars of the initial calculation. Often they will apply a safety factor or margin of extra strength to the beam design such that the beam would be able to withstand higher then expected loads. At times we may know the variability of the stress that may be applied, plus we must study and measure the full range of variability to expect in the material within the beam. Given that knowledge we can calculate the probability of the stress being sufficient to cause the beam to fail.

Electrically, designers consider the variation of power available from the grid and local power distribution system. They consider electrostatic discharge events and other common electrical power variations that the product is likely to experience. Lightning strikes either nearby or directly are immense amounts of power and very few products are designed to withstand such stress. The likelihood of a lightning strike is relatively remote and the design that can withstand such a load is very expensive, therefore, few products are deliberately designed to withstand such a load.

Product failures due to overstress occur due to design errors when estimating the expected loads; due to supply chain or manufacturing errors as described above that weaken the products ability to withstand the load; or due to a true overloading of stress, which is either expected to rarely occur or is outside the expected operating parameters of the product.

With sufficient information about the distribution of expected stresses and strengths, we can estimate the number of failures. It is much more difficult to estimate when the overstress will cause a failure though. We do not know when lightening will strike or someone drops their phone into the pool. Therefore, the approach is not to predict when it will occur, just to employ the discovery tools to estimate the margin (safety factor) within the expected operating environment. Not just to the specified operating limits, rather to the limits of what is likely to occur, including beyond the specifications.

4.2.3 Wear out

Everything fails eventually. Items wear, material is consumed, polymers breakdown, metals rust, and pn junctions decay. The intention of most designs is to create a product that provides value and delays or avoids wear out long enough to permit the value delivery.

Once the risk of failure from supply chain, manufacturing, and overstress are minimized, the remaining risk is wear out. A design that does not account for this source of failure may experience premature failure of all products placed in service. This may impact warranty claims, loss of brand image, etc. Fortunately, with a focus on the failure mechanisms discovered or known, we can reasonably estimate how long before a product will succumb to wear out failure.

There are a few common means to estimate when a product will fail. Keep in mind the amount of variation that is present in and between individual products and when, where and how they are used. The set of assumptions made around the approach for the estimate is often as important as how well the failure mechanism is known and modeled. Nominal, worst case, and Monty Carlo are methods to apply stress during a life estimate for a product.

One approach is to estimate the worst case set of stresses and apply those to the most likely to occur failure mechanisms to form a basis for the prediction. This is conservative, yet practical. Similar is the nominal conditions. This is not as conservative and rarely used. It is mentioned here, as the result between a nominal set of conditions and worst case may be significant.

Another approach is to use a random set of stress conditions drawn from the known set of stress condition distributions and apply those to life models of the dominant failure mechanisms. Repeating the selection of conditions and projecting the time to failure via appropriate life models, permits an estimate of the life distribution, not just a point estimate.

For the individual failure mechanisms or with a product that may expect a single dominant failure mechanism, focusing on the life model of the failure mechanism is appropriate. Not everything is temperature driven and modeled by the Arrhenius rate equation [3]. Thermal cycling may cause solder fatigue (Norris-Landzberg [4]), temperature and humidity may cause CMOS electromigration (Peck [5]), and there are hundreds or more models specific to a failure mechanism. The term ‘physics of failure’ implies the model of the failure mechanism is down to the physics (or chemistry) level.

It is beyond the scope of this tutorial to address all the means to characterize the time to failure behavior of a failure mechanism, yet there are many good references on the subject. Testing may focus on samples, components, subsystems, full products, and may include normal use rate and conditions or accelerated use rate and/or stress conditions. The focus is on the failure mechanism, and using a known model from literature or internal experimentation permits the team to understand how long the product is likely to last. This estimate is then compared to the reliability goal.

Besides the focus on failure mechanism models, a widespread practice for estimating reliability of electronic products is to use a parts count prediction method. There are standards that offer a listing of failure rates for components. These documents provide a means to tally up expected failure rates and predict the product failure rate relatively quickly. Telecordia SR-332 [6] is an example. Take the results of such approaches with due skepticism as they are rarely accurate and may provide a result that is over 100% incorrect. [7]

Parts count predictions like engineering judgment do play a role in estimating product life; they assist the team in making decisions. Parts count predictions also encourage reducing parts count within a product and keeping the temperature low across the components. These are good outcomes and do assist in the creation of a reliable product.

Minimizing and controlling supply chain and manufacturing sources of failure, plus designing the product to withstand the expected variation in stress and strengths involved provide a solid platform for a reliable product. The decisions made during design including material, component and assembly details will impact the time until the onset of wear out. Designed properly a product may have a long and useful life providing value. Understanding the failure mechanisms permits the team to know the likelihood of their product surviving for duration of the goal or not.

5. FRACAS

FRACAS, or bug tracking, defect tracking and similar terms relate to the process of recording issues, problems, defects, unexpected behavior or performance, testing anomalies, or product returns in order to effect product improvements. Failure happens. Recording, resolving and learning is the gift provided by a product failure.

FRACAS may be as informal as a small team discussing issues noticed the previous day to specialized database programs with hundreds of people involved. The essence is every defect or failure is captured within the system. The process usually has some form of failure analysis and triage to determine the appropriate action to take in response to the failure. Options include a product design change or adjustment, a material change, or to ignore the issue. Every failure provides information, some will require action, and often not all issues that arise will have time or resources to affect a change.

Tracking issues during the design phase helps to insure that issues identified during the design process are resolved prior to customer use. Given the limited number of prototypes generally available, every failure may indicate a relatively high failure rate once in the field, if not resolved.

Tracking issues once the product is shipped provides the necessary feedback on actual product reliability under normal operating conditions. The assumptions made during the design process are actually put to the test. If the failure rate is from the expected failure mechanisms and at the expected rate, then the work during the design process has been accurate. If not, the information provides a means to not only improve the product now, it also provides feedback to the entire process of designing a reliable product.

6. MAINTENANCE CONSIDERATIONS

Repairing a product assumes the product is repairable. Creating a product that is repairable is part of the design. Some products are not repairable simply because the repair process costs more than the value of the product. Products such as an escalator, bottling equipment or automobile have design features that make them economical to repair. The combination of the design, supply chain for spare parts and tools, and the training and execution of repairs are all part of maintainability.

There are many metrics related to the time to repair that may or may not include diagnostic time, spare part acquisition and technician travel time, along with actual hands on the equipment repair time. The time to repair along with the time to failure information is combined to provide a measure of availability. Availability is related to the concept that the equipment is ready to work when expected. Concepts of throughput, capacity and readiness are related to availability.

In the design process, the designer needs to consider access, disassembly, assembly, calibration, alignment and a host of other factors when creating a system that is repairable. For example, the oil filter on a car has standard fittings, permitting the use of existing oil filters as a replacement. The design of the system may involve tradeoffs between design features and aspects of maintainability, such as cost of spare parts and time needed to actually accomplish a repair. Cost of ownership often includes the cost of repairs and spares.

For the team maintaining equipment, the considerations include understanding the equipment failure mechanisms, the symptoms and time to failure expectations. The stocking of tools and spare parts can be expensive and minimized if the system behavior over time is understood. The team may require specialized training and certifications and that also may increase maintenance costs.

There are a couple of basic approaches to maintenance: time-based or event-based. If you change your oil every 3 months, you are using a time-based approach. If you are changing your car’s oil every 5,000 miles, then you are using an event-based approach. Both require some knowledge about the failure mechanism involved to set the triggering time or event criteria, so the maintenance is performed before either significant damage or failure occurs to the system.

Another approach is to monitor indicators of the amount of wear or damage that has occurred and repair the unit, as that unit's specific useful life is about to fail. For example, periodically testing an oil sample may reveal when the oil is about to become ineffective as a lubricant. Monitoring and maintenance can be very sophisticated or as simple as having a brake wear indicator that causes a squealing sound. Prognostic health management is a relatively new field focused on measurement techniques that, like the wear indicator in brake pads, assists the maintenance team in maximizing the useful life of a product and effecting repairs and maintenance only as needed to prevent failure.

7. VALUE

The various tasks and activities commonly associated with reliability and maintainability are not accomplished without purpose. They add value to making decisions, provide valuable direction and feedback. These tasks help to avoid expensive mistakes or excessive repairs. These tasks guide designs to become more reliable and cost effective. They are done to add value.

Conducting HALT on a prototype that the design team ignores is a HALT of little value. A prediction done only to meet the contract requirements and not reviewed and acted upon by the design team is of little value. Conducting a set of 'reliability tests' that are not related to failure mechanisms or use conditions again is of little value.

A simple question to ask when planning or starting any reliability or maintainability task is: "How will this information be used?" This is like asking for information about the audience of a presentation. If the task does not produce information of value then it is appropriate to not spend time and resources on said task.

8. CONCLUSIONS

Reliability and Maintainability Engineering are challenging and rewarding endeavors. Managing to bring a product to market that provides a valuable service over it lifetime is difficult. The tools and resources of R & M engineering provide a means to efficiently achieve the reliability and maintainability goals.

The basic outline used in this tutorial provides a guide to establishing an effective means to manage reliability and maintainability.

- Set a goal
- Articulate the goal clearly throughout the process and organization
- Discover the salient failure mechanisms
- Minimize supply chain, manufacturing and overstress failures
- Estimate the product's life
- Track and eliminate failures

The implementation will be different for every organization. Yet even this simple outline does permit the entire team to make decisions leading to a reliable or available product. Focus on failure mechanisms and obtaining a solid understanding of the dominant failure mechanisms behavior over the range of use conditions. And, finally, only do the tasks and activities that add value to the organization.

9. REFERENCES

[1] Petroski, H. (1994). *Design Paradigms: Case historyes of error and judgement in engineering*. Cambridge, Cambridge University Press.

[2] Holmes, Oliver Wendell, *The One Hoss Shay*, New York, Houghton, Mifflin and Company, 1891.

[3] Ireson, William Grant, Clyde F Coombs, and Richard Y Moss. *Handbook of Reliability Engineering and Management*. New York: McGraw Hill, 1995, p. 12.2.

[4] Norris, K C, and A H Landzberg. "Reliability of Controlled Collapse Interconnections." *IBM Journal of Research and Development* 13, no. 3 (1969): 266-271.

[5] Hallberg, Ö, and D S Peck. "Recent Humidity Accelerations, a Base for Testing Standards." *Quality and Reliability Engineering International* 7, no. 3 (1991): 169-180.

[6] Reliability Prediction Procedure for Electronic Equipment, SR-332, Issue 3. Telcordia, January 2011.

[7] Jones, J, and J Hayes. "A Comparison of Electronic-reliability Prediction Models." I*EEE Transactions on Reliability*, Vol. 48, no. 2 (1999): 127-134.