

What is Equipment Reliability and How Do You Get It?

By Mike Sondalini and Howard Witt

Abstract

High equipment reliability is a choice and not an accident of fortune. To a great extent you can choose how long you want between equipment failures. You can deliver high equipment reliability by ensuring the chance of incidents that cause failures of equipment parts are low. The secret to remarkably long and trouble-free equipment lives is to keep parts and components at low stress, within good local environmental conditions, so there is little risk they are unable to handle their design duty. If there is nothing to cause a failure, the failure will not happen and your equipment continues in service at full capacity and full availability.

For maintenance practitioners, it is the use of precision maintenance backed-up with condition monitoring applied as a tool to prove maintenance work is done precisely, that provides the foundation for exceptional equipment reliability.

Keywords: equipment reliability, downtime, failure avoidance, precision maintenance, plant availability, equipment risk reduction, condition monitoring, reliability growth, hidden factory

Without getting into the mathematics, equipment reliability is a measure of the odds that an item of equipment will last long enough to do its duty. It is a measure of the chance of remaining in-service to a point in time.

You measure the reliability of equipment by its trouble-free time. If it is meant to last for 10,000 hours (about 14 months of continuous operation), and does last that long, it is 100% reliable to 10,000 hours. But if after 10,000 hours there is an occasional failure, the reliability beyond 10,000 hours is less than 100%. When we talk about reliability, we must also say what time period is involved. When equipment operates at duty capacity for as long as expected, it is considered reliable. When the period between out-of-service episodes is too short, it is unreliable.

Equipment reliability needs to be seen as more than just a chance time span. It is about building great businesses that are world-class performers. High-reliability organizations expect equipment to last a long time and are unhappy when it does not. Not only are they unhappy, but they take effective measures to learn and improve from the failures.

You have to deeply want the production and profit benefits equipment reliability brings before you will do what is necessary to get it. If you want plant and equipment to operate trouble-free for a long time, you must do those activities that cause reliability, and do them well enough to deliver reliability. You get great equipment reliability when you act to control happenstance across the life-cycle and replace it with masterly precision. If you want high reliability with low cost, put into place the necessary engineering, purchasing, storage, operating and maintenance regimes and practices that deliver the reliability and life-cycle costs you want.

You may have limited opportunity to influence any of these regimes on your plant. In that case, start by looking to get more value from the equipment you have. Improve those plant items with lower than desired reliability - the 'bad-actors'. You can do that by using operating practices that reduce equipment risk and by improving machinery health. Use cross-functional teams of operators, trades people, condition monitoring technicians and engineers tasked to understand the causes of a problem and to eliminate or reduce them. Spot the onset of problems and take pre-

emptive action. Have contingency plans to mitigate consequences. Operate and maintain plant to precision standards. These risk reduction practices will deliver higher plant reliability.

Measuring the Impact of Equipment Reliability on a Business

The standard definition of reliability leaves much unsaid about the effects of equipment failure on businesses and people. You know when you have unreliable plant and equipment because people are angry that it fails so often. In companies with equipment reliability problems people are busy ‘doing’, often repairing failures over and over again. It never ends, and you go home each day knowing there will be more troubles tomorrow. You also know when you have reliable equipment because it performs as its design intended without failures. The business likely makes good profits with low operating costs controlled to a narrow, known range. You have the time to do your work well. A place with reliable equipment is a happy and safe place.

Measuring equipment reliability is important if you want to improve it. Reliability is measured as the average time between failures, known as ‘Mean Time Between Failure’ (MTBF). One drawback with only measuring time is that there is no indication of the value of that level of reliability. If you do not know what reliability is worth, you may spend lots of money on small improvements that have little impact on profitability. Or worst still, not spend enough money on highly profitable improvements. Reliability measured only by expected time in service is a poor business indicator. As Benjamin Franklin, the publisher of *Poor Richard's Almanack 1733*, said, “Time is money.” Reliability must also be measured by the money made or lost for the business. A measure of improving reliability needs to show improving profitability.

For a company that measures production by weight a standard profitability indicator that reflects reliability is:

$$\text{Unit Cost of Production } (\$/T) = \frac{\text{Operating Costs in the Period } (\$)}{\text{Total Saleable Throughput (Tonnage)}}$$

Figure 1 shows you what these measures indicate.

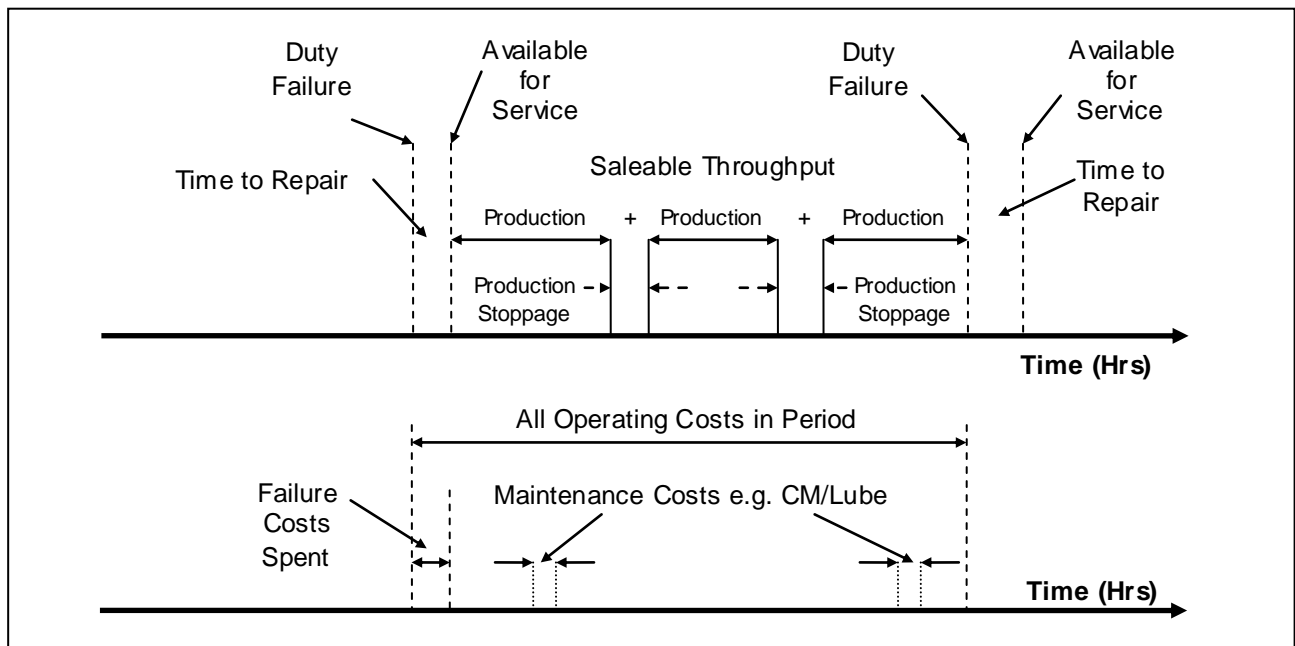


Figure 1 – Reliability ought to Measure Time and Operating Cost

Let's take a simple example of a plant intended to operate for another 10 years (87,600 Hr) producing 1,000 T/Hr which currently has only 90% Availability due to equipment breakdowns. If it continues to be run with unimproved reliability, it will lose 876 hours of potential production per year, or over 5 weeks lost yearly, which represents about 10% of annual production.

For a plant where little attention has been paid to reliability improvement it is possible to double MTBF (i.e. halve the number of breakdowns) using failure elimination methods such as precision maintenance, and machinery improvement projects aimed at defect removal. With focused efforts on better planning, improved maintainability and skills upgrading, the MTTR (Mean Time To Repair) for maintenance work could be halved. Of the 5 plus weeks lost we could expect to recover 3.8 weeks with the above strategy. That is an additional 507 hours per year and represents an increase from 90% to above 97% Availability.

Under this simple example the effect on the Unit Cost of Production is to lift the Saleable Throughput by 7% while driving costs down. In halving the breakdowns the direct maintenance component of the Operating Cost (such as parts, manpower, and supervision) falls markedly. To this must be added the huge reductions in indirect costs, plus the opportunity and knock-on costs no longer lost. Together these three cost categories can easily be ten times the value of the direct maintenance saved. In such a case, for every \$1,000 dollars of direct maintenance cost reduction through improved reliability, you also gain back \$10,000 in additional profit previously lost. Improved reliability causes throughput to rise, lowers direct production cost and avoids consequential business-wide losses. On a graph of the Unit Cost of Production over several periods you would happily see a falling trend as the reliability growth strategy changes the place.

Unit Cost of Production provides a means to show the effect of reliability in real money and production. Higher reliability reduces the unit cost in three ways - you get more time to make product, you get more product out, you keep the earnings the business would have lost from failures. As reliability improves there is more production time, throughput and 'added' profit. That is why every company wants more reliable equipment; there is a lot of money to be made from plant reliability improvement.

Of course there is a price to make the needed changes. Making plant access refinements, doing equipment design upgrades and training tradesmen and operators in higher and better skills is added expense and takes time. But when you can show people returns on investment measured in the hundreds of percent, you will get the management support and capital you need.

We now have two measures to show the impact of equipment reliability improvements – increased 'Hours between Duty Failures' and reduced 'Unit Cost of Production' over a time period. These are all-encompassing measures that include the effect of everything impacting the equipment's reliability. To actually improve equipment reliability we need to rip those measures apart and look inside them to find what caused those results.

It is best if you know the costs spent on an item of equipment, exactly what they were spent on, and what they were spent to do. If you do not have good records of equipment problems, with their associated costs and changed conditions over the equipment life, it will probably be difficult identifying suitable maintenance actions and preparing justifications to fix them. To get the detail needed to fully analyse the cause of your reliability problems requires tracking the life and cost of your equipment down to the individual part level.

The secret to successful equipment reliability lies within the lives of your equipment's parts.

Equipment Reliability Depends on the Reliability of Parts and Components

Equipment is made of parts and components combined in assemblies that work together to allow it to operate. Figure 2 shows how we combine parts to form a shaft bearing assembly. It shows a bearing in an electric motor housing carrying a shaft; a typical situation in many industrial machines. There are 14 parts in the assembly. The 14th item is the lubricant.




Figure 2 –Series Arrangement of Parts Allow a Shaft to Turn in a Electric Motor Bearing Housing

When you look closely at how the assembly is built, you find that it is configured such that parts work with others in a sequence. Figure 3 identifies a portion of the sequence of parts that allow the shaft to turn in the bearing. Notice that they are organised in a series arrangement.




Figure 3 – The Series Arrangements of Parts that Allow the Shaft to Turn in a Bearing Housing

All industrial equipment is built as a series arrangement of parts and components working together to perform the required duty. Once you have a series arrangement of working parts, the series reliability depends on each one working properly. A simple example shows you what happens.

Figure 4 is as Figure 3, except with a failed part. Without correct lubrication the series connection has been lost and the assembly cannot survive in-service. If this assembly were in a piece of equipment, the equipment would be failed. A sequence arrangement of parts only requires one item in the series to fail and the whole assembly fails. When the assembly fails, the equipment stops.




Figure 4 – The Reliability of a Series Arrangement Depends on the Reliability of its Components

The intended message in Figures 3 and 4 is that the reliability of a piece of equipment is totally dependent on the reliability of its individual parts. If one part fails the machine fails. If one part is in a bad condition the entire equipment is at ever increasing risk of breakdown. We can calculate the reliability of a series system mathematically. In Figure 5, each part has its own reliability – ‘R’.




Figure 5 – Every Part in a Machine has its Own Reliability

Equation 1 is used to calculate the reliability of items in series. It is the multiplication of the reliabilities of the individual items in the series. It makes more sense if we use numbers. For the sake of the example, say the reliability of each item is 99% (usually written as 0.99), which means 99 parts out of 100 are failure-free for as long as they should be.

$$R_{\text{series}} = R_1 \times R_2 \times R_3 \times \dots \times R_n \quad \text{Eq 1}$$

The series reliability for the eight parts is:

$$R_{\text{series}} = 0.99 \times 0.99 \times 0.99 \times 0.99 \times 0.99 \times 0.99 \times 0.99 \times 0.99 = (0.99)^8 = 0.92 \text{ (or 92\%)}$$

Once parts with reliability of 0.99 are in a series assembly of eight items, the reliability of the series falls to 0.92. Only 92% of the assemblies will last to the time they should. As the number of parts in a series grows longer, its reliability declines because there are more things to go wrong. Watch what happens to reliability when a part has a serious problem. If the lubricant has a reliability of 0.5, where half of the time it never reaches its expected life, the equation becomes:

$$R_{\text{series}} = 0.99 \times 0.99 \times 0.99 \times 0.99 \times 0.99 \times 0.5 \times 0.99 \times 0.99 = 0.47 \text{ (or 47\%)}$$

Now only 47% of the assemblies last their full life. It only takes one part in a series arrangement to be of low reliability and it brings the entire machine down to below that reliability. We have now arrived at a most important principle in equipment reliability:

An assembly of parts (i.e. a machine) can never be more reliable than its least reliable part!

You can never improve an equipment item’s reliability more than its least reliable part. Understanding this reliability principle is important for everything you do in the field of maintenance and reliability improvement. If you want to improve your machines’ and equipment reliability, you first must ensure each of their parts are even more highly reliable.

When Machines and Equipment Fail We Replace Parts

When we do Preventive Maintenance or Breakdown Maintenance, we replace parts and/or lubricant in a machine and then put the equipment back into service. The new parts start their life, while the parts not replaced continue theirs. There is also the very real possibility that parts which were minimally stressed before the invasive maintenance become stressed due to poor maintenance, and even that some of the new parts installed are stressed during assembly. Now within the machine there are old parts still in good health, parts that have accumulated stress and approaching end-of-life, distressed parts ready to fail from accumulated overloads, and new parts starting into service with their inherent design limitations. What is the reliability of the whole machine now?

We know that equipment reliability depends on individual part reliability. The distressed parts have a very poor reliability (likely to fail soon); while the new parts should have much higher reliability, (likely to fail sometime in the future). Overall, the equipment is no more reliable than the most distressed part. What could you do right now to improve the reliability of the distressed part?

You could stop the equipment and replace that part with new. The Operations Group would be very unhappy to learn that the equipment again needs to stop. Also, you must know which parts are in distress, else you may replace the wrong ones and the equipment will still fail soon. There is another thing you can do – reduce the chance of overstressing the part. If the chance of excess stress is substantially reduced, the distressed part has a greater prospect of lasting longer. Lowering the stress on machine parts greatly improves the odds for higher equipment reliability.

This can be as simple as improved housekeeping, such as keeping breathers clear of dust to prevent lubricant contamination and cleaning rubbish/rags and dust/dirt build-up off electric motors, bearing housings and gearboxes to improve heat loss. You can also use CM to monitor the stress from ‘rough operation’ induced by poor operating practices and bring the likely implications on production to the attention of Operations and Plant Managers.

Measuring the Rate of Equipment Failure

Parts working together form a series system we call machines. When any working part fails, the machine fails. You can draw the failure rate curve for a machine from the rate of its parts’ failure.

Figure 6 shows the life of an imaginary machine with three working parts. An example could be the bearing, seal and shaft in a bearing housing of a centrifugal pump or an electric motor. The bottom chart shows the ‘Hazard Rate’ curve of the individual ‘green’, ‘blue’ and ‘red’ parts. Such curves represent numbers of parts expected to fail in a period from particular circumstances. The ‘blue’ part has wear-out characteristics and is replaced on breakdown. The ‘red’ part has an infant mortality characteristic. Sometime it fails early, other times it is later. The ‘green’ part is characterised by a life of random failures that can happen at any time. When a part fails (shown by an ‘explosion’), the machine also fails.

The top chart reflects the whole machine’s rate of failure, which is the sum of its parts’ failure rates. The more often they fail, the higher the machine’s failure rate. When parts do not fail, the rate falls. The machine’s failure rate curve is called ROCOF - Rate of Occurrence of Failure. The ROCOF is representative of the reliability of the machine design, the quality of manufacture, the precision of its installation, its production abuse, the purchasing and storage quality control, along with the standard of maintenance and workmanship care.

Figure 6 tells you what to do to prevent equipment failures – you must first stop parts’ failures.

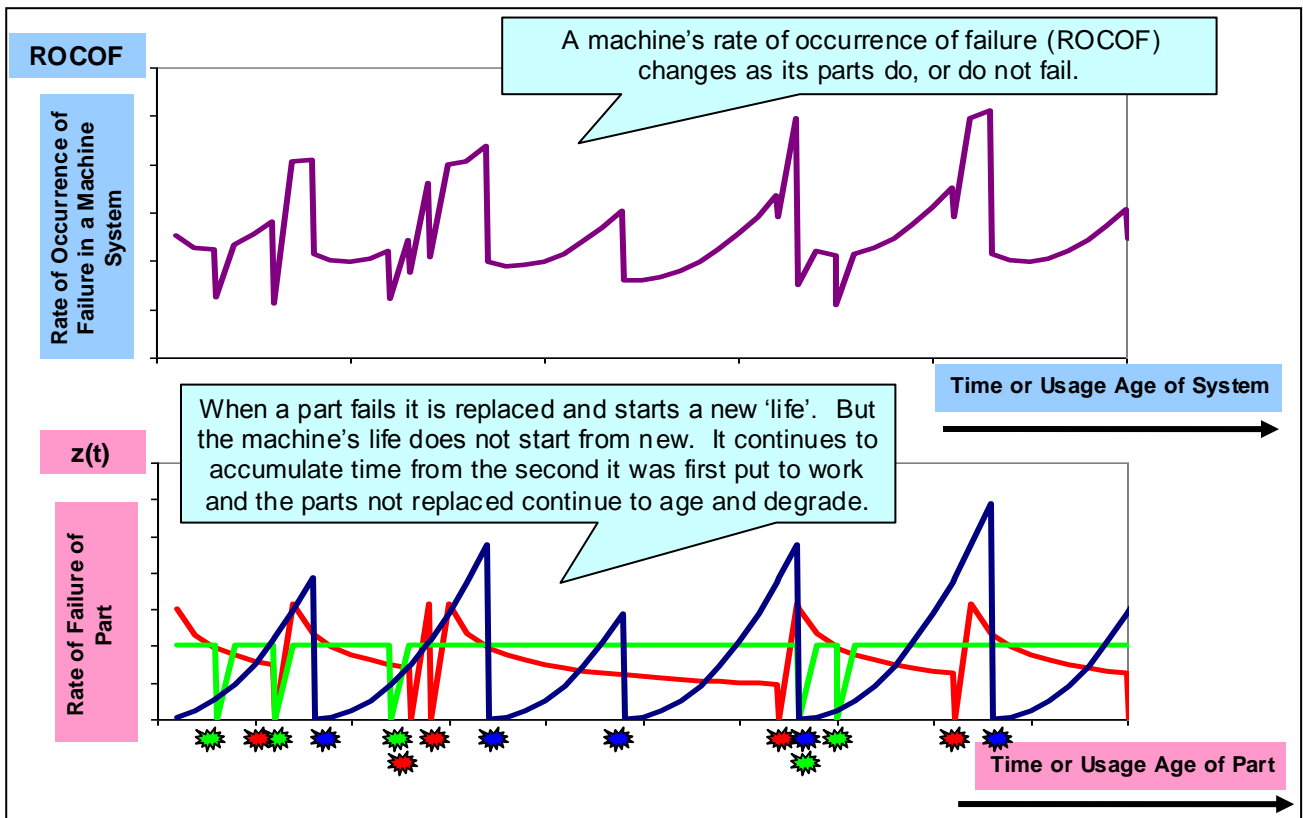


Figure 6 –Machine Failures are the Accumulated Effect of Its Parts' Failures

The ROCOF curve for a machine reflects what happens to its parts, and moves up and down as parts fail. When we take the parts' failure history of many identical machines we get the mean, or 'steady average' ROCOF shown in Figure 7. Figure 7 also lists many of the reasons why equipment and machines fail during their lives.

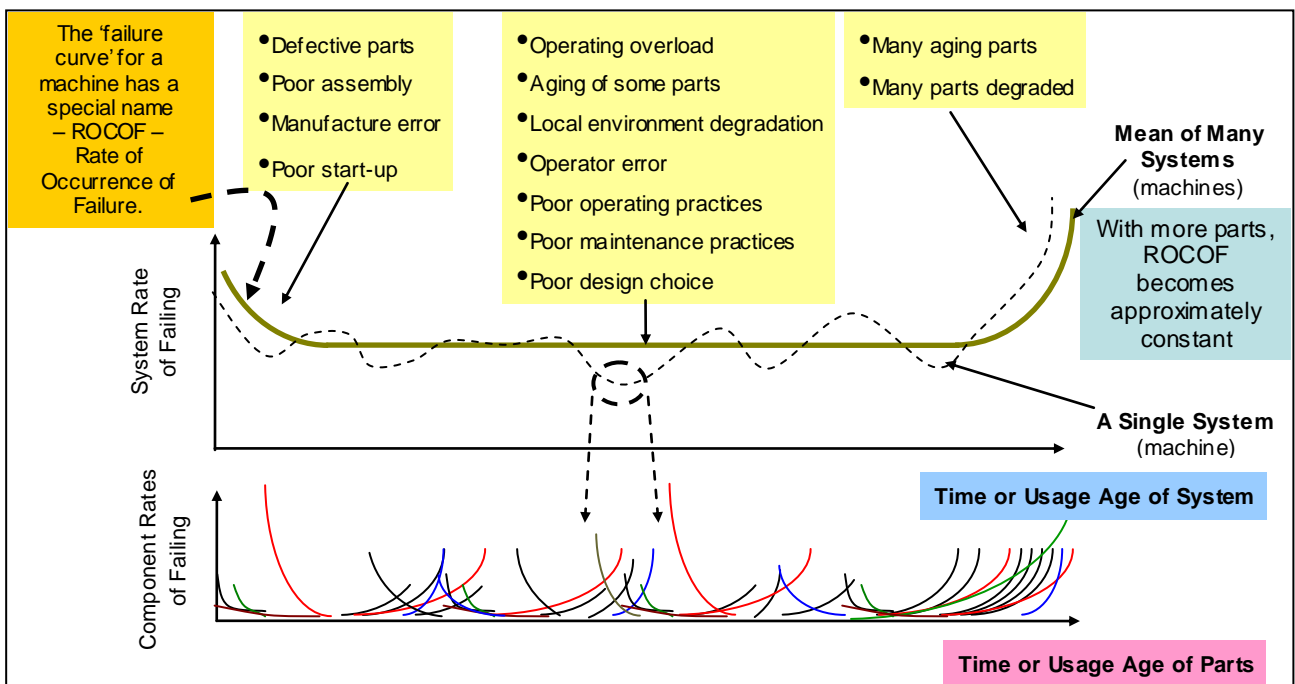


Figure 7 – ROCOF Curve for a Machine (i.e. a System of Parts) Design

First parts fail and then machines stop. The solution to equipment reliability is to improve parts' lifetime reliability. The limiting (maximum) reliability of parts is set by their design. Once a part is

in a machine, we are limited to its characteristic performance. At best, it will behave as its design allows. This is the limit to how much reliability you can get from a part without redesign.

With ideal maintenance and no operational over-stressing we can achieve the design limits and get the designed equipment reliability. We can even do better than the design limit, and lower the equipment failure rate, if we de-rate equipment and use lower loads or more benign environments on the parts.

In practice industrial equipment failures are typically more frequent than expected for their design limits. This higher failure rate can be caused through wrong installation, bad operating or housekeeping practice, bad supply chain and stores management or less than ideal maintenance. Fortunately, the induced failure rate of machines is highly malleable depending on your choice of the applied maintenance policies, the operating policies, purchasing and stores practices, installation accuracy and the assembly precision of the machine's parts.

You may have limits on purchase options and you may be stuck with the equipment you have, but correct care, operation and maintenance of what you have is totally within the control of every business. Figure 8 shows what happens to the ROCOF when parts' failures are removed.

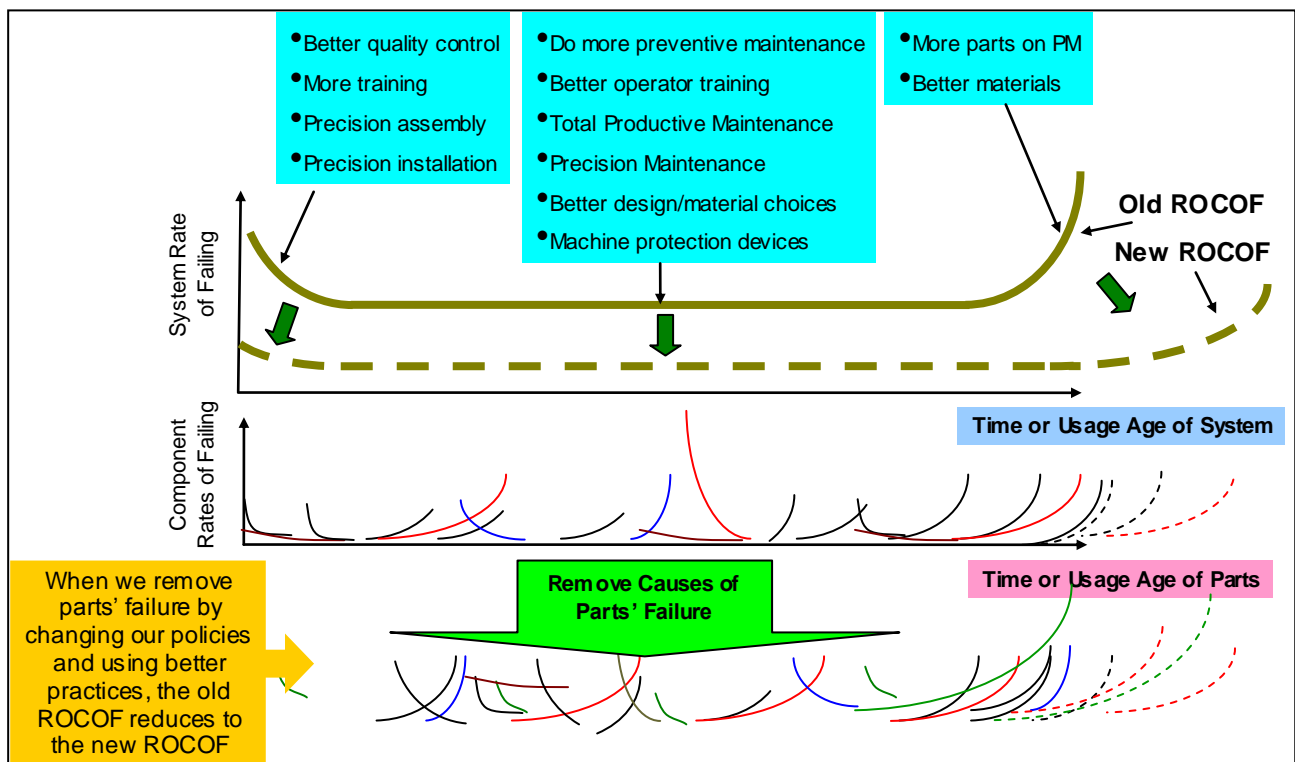


Figure 8 – Improving the Reliability of Machines and Equipment

Stop the Risk of Excessive Stress and You Stop Equipment Failures

How long will your equipment last before the next failure? You cannot possibly know with certainty because it depends on the chance its parts will survive to a point in time. The best you can offer is a guess. I bet that you did not know you are gambling the future of your business when you work in Engineering, Operations or Maintenance. Reliability worsens and plant availability falls if you do not understand the game you are playing. But to turn the odds in your favour you only need to control what stresses you permit your equipment parts to experience.

Risk is a combination of the consequence and likelihood of specified outcomes. It is usually calculated by the following equation.

$$\text{Risk (\$/yr)} = \text{Cost of Failure (\$)} \times \text{Chance of an Occurrence (/yr)} \quad \text{Eq. 2}$$

Thus risk is a quantified guess/estimate of how much attention should be given to a set of potential outcomes/scenarios (e.g. failures). Regardless of values estimated by Equation 2, you will be hit with a cost every time the failure happens. After a very bad event people may or may not be understanding of your view that you paid the issue as much attention as was appropriate for the risk.

Restating: to reduce risk we must first reduce the chance of a bad event, because once it happens you will pay the full price of that occurrence. If you do not want to incur the associated cost penalty you must not just guess, but be confident, that the chance of a failure event is truly very low. Conversely, if you want beneficial events to happen (e.g. like winning a lottery / having well lubricated parts), you must provide for the chance of the good occurrence (i.e. buy a lottery ticket / rigorously follow good lubrication management practices).

Variation and the Need for Accuracy

We can greatly reduce the risk of bad events through our choice of actions and our degree of diligence in doing them. Laser aligning shafts is far more accurate compared to using a straight edge. Using laser alignment greatly improves the chance that the shafts will run more concentrically. But the risk to the equipment remains high if the alignment is not done correctly and the final operating shaft positions are out of tolerance. Our policy decision is good – use laser alignment – but if the implementation is poor the machine still fails early from higher stresses induced in the working parts than their design limits. Using good policy, like laser alignment, without also controlling the quality of the execution does not guarantee improved reliability. Variation in performing the alignment must also be controlled.

Figure 9 shows a curve representing variability in a parameter and associated outcomes (a normal distribution, though other shapes are encountered) following some action such as maintenance. An example where this curve applies would be the dimensions of a shaft or bearing where both undersize and oversize lead to bad outcomes. Another example is over torque or under torque on a bolt, where either could lead to failure.

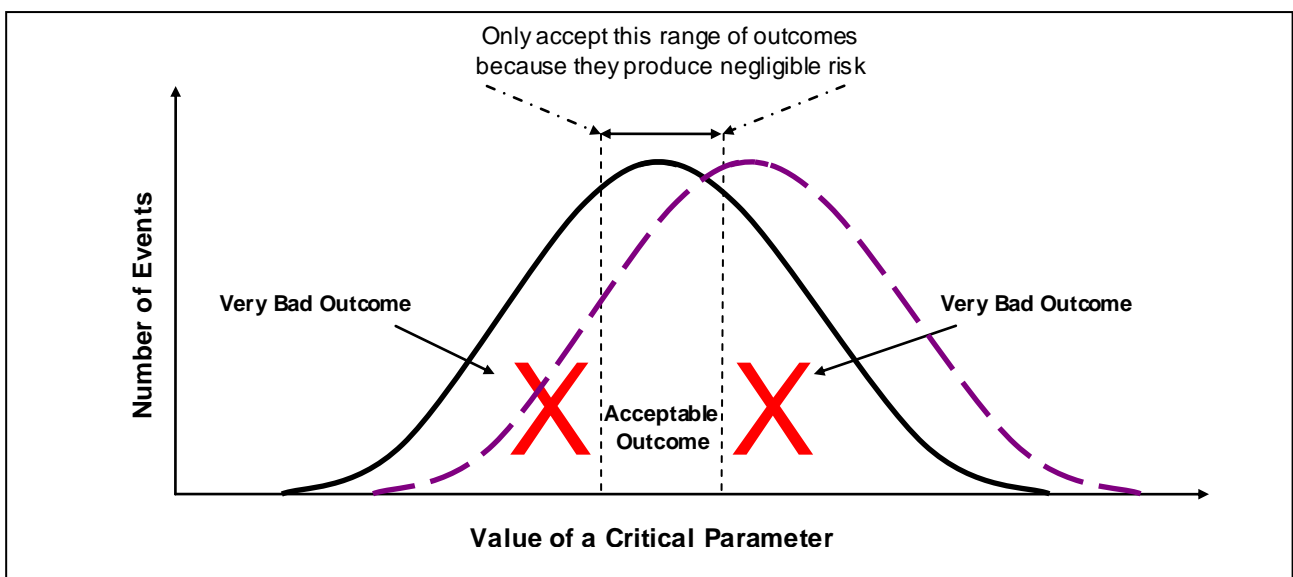


Figure 9 – Variability in Outcomes

With the solid black curve the individual or team doing the action intend to aim at the acceptable range of values, but lack of precision and poor control over quality lead to excessive variability in the parameter. Generally they get it right, but not always. Quite often they end-up with the parameter value either too high or too low. With the purple dashed line they are no less precise but get more bad outcomes because their target value for the parameter is too high.

Such curves help us to understand the extent of the risk placed on our equipment parts by our activities. If we know which outcomes cause our parts to fail, and which extend their life, we will want to ensure we do those occurrences that lead to good lifetime reliability and prevent those that cause increased chance of failure.

To reduce the chance of equipment failure you need to prevent situations where its parts are overstressed, or experience fatigue, or suffer contamination. Focus your efforts on removing the causes of bad-chance to equipment parts. If there is no bad event a part continues its natural life unchanged. With the chance of bad occurrences reduced; equipment risk falls, reliability rises and money is not spent.

We want to create the situation shown in Figure 10, where controls are applied intentionally to reduce the range of outcomes to those that are beneficial. Such controls give more accuracy and precision to achieving the parameter's outcome. This exactness leads to reduced chance of failure. The change in performance of the individual or team represented by the move from Figure 9 to that of Figure 10 indicates use of appropriate discipline, skill and motivation to work to the target range.




Figure 10 – Control the Chance of an Equipment Failure Event

Control the Chance of Overstress and Fatigue Occurrences and You Control Reliability

Figure 11 identifies many strategies available to control the risk placed on equipment parts. Those from the left-hand column reduce chance of failure. Those on the right only reduce cost of failure, but the chance remains unchanged. Notice in Figure 11 that Condition Monitoring (CM), as it is usually done, does not reduce the chance of failure. It only spots impending failure and lets you turn the repair work into a planned job instead of a breakdown. The failure has already been initiated, and if it is not addressed in time your equipment breaks down.

Figure 12 shows the effect on availability depending on your business focus in using CM. When failures badly affect production equipment CM is used to observe tell-tale performance parameters and react to their rate of deterioration. Used this way CM improves reliability by extending operating life right up to near-failure. But condition monitoring can also be used as a tool to stop failures, and not only a tool to spot failures. Depending only on the time when you start condition monitoring it becomes a tool to reduce the chance of parts failing.

When CM is used to also gather information on the chance of failure, it becomes a tool to optimize plant availability. With it you confirm that variation has been controlled to the precision outcomes you want for parts installation, lubrication and workmanship. You use it to prove the starting quality of work performed and the starting machine condition. You do this proof-testing as part of the maintenance or installation work.

For example, check for poor alignment or soft-foot with vibration analysis when re-commissioning. If vibration levels are too high after assembly, identify the cause and rectify it before letting the equipment go into service. Test for poor electrical connections (hot spots) with a thermography camera before handing the equipment back for operation. Take an oil sample when the equipment is running at operating temperature and get the start-up wear particle count. In this way, CM is used to prove the equipment has been set-up with the best chance of achieving high reliability. It is the evidence you need to prove to people that risk is truly minimised. As a direct consequence you will get lower operating costs and gain plant availability.

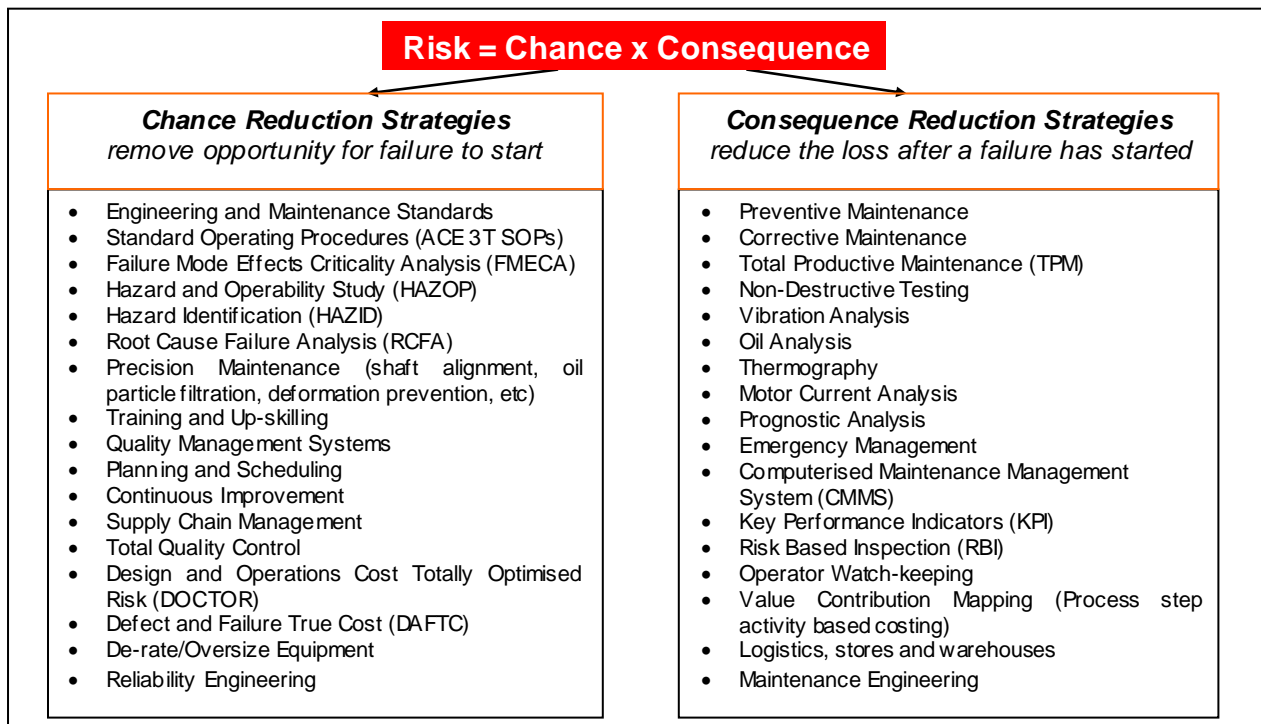


Figure 11 – Various Risk Management Methods

Condition Monitoring used to spot ‘bad-chance’ starting checks if plant and equipment is set-up in a state to deliver high reliability. By using CM in this way you optimize availability, because for a small maintenance cost to do the CM, you reduce the chance of equipment failure from undetected defects. Maintenance done this way makes money for the company by stopping the business-wide costs that would have resulted from the failure. By guaranteeing high reliability at the start, your failure avoidance delivers drastically lower operating costs. This effect on operating cost of using CM as an optimising tool is shown in Figure 13.

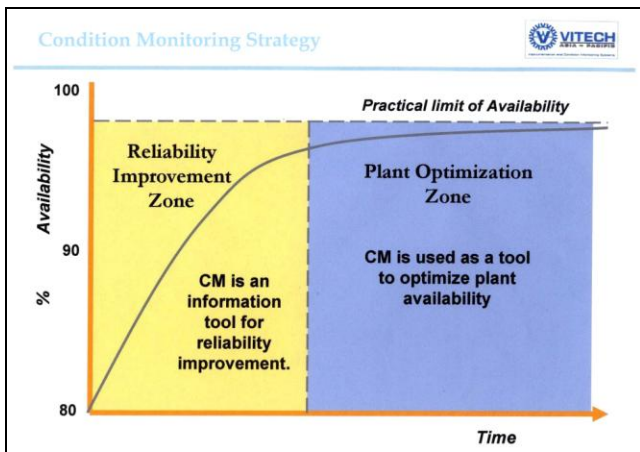


Figure 12 - Using Condition Monitoring to Control Risk of Low Reliability




Figure 13 - Condition Monitoring can be Used to Optimize Availability

Through persistent use of Condition Monitoring to measure work quality and achieve start-up precision you will build knowledge and improve workmanship standards that bestow plant and equipment in excellent condition. You will deliver low operating risk to parts and the maximum chance of a highly reliable and long, trouble-free equipment lifetime.

Summary and Conclusion

It is appropriate to draw together the key issues on equipment reliability that is in the article.

1. The reliability of a part is the chance it will survive in-service for a required length of time.
2. The level of equipment reliability has immediate influence on your business profitability through the direct and business-wide consequential impacts of failure.
3. Reliability is malleable by the design, selection, manufacturing, storage, operating and maintenance standards you allow.
4. Plant and equipment are a series arrangement of individual parts. Such configurations carry high equipment risk, since any cause of a working part's failure can stop the equipment.
5. The reliability of equipment depends on the reliability of its individual parts. For high equipment reliability each part must have far higher reliability than needed for the equipment.
6. If you reduce the chance of parts' failure by any suitable means (many of which are in the control of an operation's management, operators and maintainers) you lower risk and increase equipment reliability.
7. Prove that the precision standards needed for high parts reliability are present at start-up.

Getting high equipment reliability is mostly within the power of every business. You improve equipment reliability, and hence business profit, by choosing the policies, using the methods, and adopting the standards that reduce the chance of bad events happening or that increase the chance of beneficial outcomes.

You can use condition monitoring as a tool to detect the onset of failure. But you get far greater worth from it, if you also use it to ensure that the high quality work and precision standards which produce long lifetime reliability are present for your machine parts at the start of their lives.

Our best regards to you,

Mike Sondalini and Howard Witt
www.lifetime-reliability.com