

Lessons Learned: Tragedy and Triumph

April 1, 2006

by Reid Willis

Washington Chapter, Society of Reliability Engineers

Reid Willis is a consulting engineer with 30 years of experience in system reliability and maintainability analysis, design, testing, and program planning.

In R&M engineering, just as in other human endeavors, we learn from our mistakes and our successes. The objective here is to improve our success rate, as a community of reliability engineers, by sharing the hard lessons we have learned in our profession.

Some of these tales are happy and some are sad. They are offered in the hope that there is something to be learned from both victory and defeat. We solicit suggestions from readers; please e-mail your comments and contributions to reidwillis@juno.com. Your experience could help others in the discipline. Published by Barringer & Associates, Inc. in 2006 with permission by Reid Willis.

[Note from HPB: Reid has taken complex subjects and reduce the complexity to a one or two sentence lesson learned: The lesson is elegantly clear, concise, and too the point!]

Topic 1, Program Planning

For every problem there is a solution that is simple, clear, easy and wrong.
– H. L. Mencken

1.1 Program Initiation

In a seminar on the impact of system R&M characteristics on support requirements, the manager from an aircraft manufacturer noted that he had found that 70% of their products' life cycle problems arose from the concept design. A manager from another major manufacturer agreed that most of their operational support problems were rooted in the initial design.

At the same seminar, a panel found that none of its members from government and industry knew what an R&M program was, what was in it, or how it was managed.

Lesson learned:

It is essential for the achievement of long-term system R&M performance that a qualified R&M program manager be empowered and an R&M program be established and funded at the outset of the overall project.

Topic 2, Specifications

Be careful what you specify. You might get it.

2.1 That's Obvious

A company developed specifications for a liquid cooling chamber. The chamber's temperature range, rate-of-change and tolerance were all specified. The test group recommended adding a specification that the chamber meet all requirements with one set of adjustments, but management said that was obvious.

The manufacturer delivered the chamber and demonstrated satisfactory performance. However, different adjustments were required to achieve specs at high, low and ambient temperatures. Additional adjustments were needed to demonstrate the rate-of-change requirement. When it was pointed out that one set of adjustments should meet all requirements, the manufacturer asked for more money and more time on the basis that this was a change of scope.

Lesson learned:

The first rule of writing specifications is to state the obvious.

2.2 A Slight Misunderstanding

An airline contracted to buy a number of aircraft from the manufacturer. When the first planes were delivered they failed the airline's acceptance test. The problem turned out to be a difference in interpretation of the contract specs. The manufacturer had to bear the cost of a design change and the airline's plans for using the aircraft were delayed.

In a review the manufacturer found a history of problems that arose during acceptance testing and were due to misunderstandings concerning contract specifications. They adopted a policy of appointing a team of operators and engineers, who had not been involved in the project, to review all specifications and acceptance test plans together with customer engineers, before the contract was signed.

Lesson learned:

All misunderstandings concerning contract specifications and acceptance test procedures will eventually be resolved. The best time to do that is before the contract is signed, not after the product is delivered.

2.3 System Requirements vs. Specifications

A government study of new-technology projects found that contractual reliability specifications were often unrelated to mission reliability requirements.

Lesson learned:

When the company is preparing contract reliability specs, get involved.

Topic 3, Data

There are only two kinds of reliability engineers: those who say “Data is the problem” and those who say “Data are the problem.”

3.1 Planning Task Data Acquisition

The Coast Guard planned to put a class of cutters through a major overhaul. Two reliability engineers were assigned to recommend reliability improvements. They chose mission operational availability as the figure of merit and mission simulation as the analytical approach.

They foresaw that acquiring failure and repair data on the ships' equipment would be a major problem. Military equipment data bases did not apply to most Coast Guard systems and the Coast Guard did not keep maintenance histories. They planned a data strategy in four steps.

1. Expend half the planned data collection time and money to get the best data they could obtain with those resources.
2. Fill in the blanks with worst-case parametric data.
3. Run a trial simulation to identify those equipments whose rates could significantly affect mission A_0 .
4. Expend the remaining half of the data collection resources to resolve the critical equipment values.

Their strategy was successful. It allowed them to concentrate on the data elements that had the greatest effect on the ship availability.

Lesson learned:

Plan the acquisition of needed data at the outset of the task. Adopt a strategy that will identify the most critical data elements and focus the available task resources on those values.

3.2 Which Data to Use?

A reliability team was creating a computer model of an existing aircraft carrier for mission simulation. The team screened maintenance records to develop equipment MTBF and MTTR values. The client was delighted to see “real fleet data,” but the team was concerned that the machinery operating hours and active repair times were rough estimates. The team manager preferred operations and maintenance supervisor judgment over any other data category. He hired a retiring chief petty officer with extensive carrier experience to review the figures. The chief spent several days telephoning old shipmates and marking-up the data sheets. For example he found it was standard practice to shut down one engine during transit, significantly reducing engine operating hours. In another example the forced draft blowers were mounted in the overhead in that ship, greatly increasing the required repair time. It is not clear whether the analysis results were different but the data values certainly were.

Lesson learned:

There is no hard-and-fast ranking of preferred data sources. That would be too easy.

3.3 Controlling Task Data

The contract for procurement of a new kind of military vehicle required the performance of R&M tasks at each phase of design, prototyping and production. The tasks included reliability and maintainability prediction and allocation, failure modes effects and criticality analysis, and mission simulation. The producer’s logistics manager was assigned responsibility for submitting the initial R&M analyses and periodic updates to the government. The corporation granted him his pick of newly-hired engineers to do the analyses and an R&M consultant to show them how.

The manager understood the importance of data control. His first move was to assign that responsibility to his strongest engineer. The engineer established a computer spread sheet and published it on the project LAN. The spread sheet contained a level-of-indenture listing (overall system, major subsystems, etc.) of the vehicle design down to the equipment level, the equipment MTBF and MTTR data, and a key to the data source. The purposes of the spread sheet were to:

- Publish the system configuration and equipment failure and repair rates to be used for all analyses.

- Clearly depict any shortfalls in the engineering divisions' responsibilities for providing and updating data.
- Notify all engineers and R&M analysts of design, equipment, and data changes.
- Provide a basis for quick-response R&M support of trade studies.
- Synchronize the set of R&M predictions and analyses submitted at each due date.

When problems arose concerning analytical procedures, conclusions and recommendations, the R&M manager and consultant were free to solve them without major data fumbles. Almost half of the submitted recommendations resulted in R&M design improvements, which is a high batting average in this business.

The logistics manager later applied the same kind of positive data control to R&M tasks in smaller projects and it proved equally effective there.

Lesson learned:

At the outset of an R&M program, establish a spread sheet of the system configuration, R&M data, and data sources. Use the spread sheet as a tool for controlling data collection, data consistency, and analysis updates.

3.4 Planning Ahead

In order to maximize aircraft operational availability, airlines periodically replace certain kinds of equipment on a schedule before they are expected to fail. They keep "rotatable pools" on hand so they can quickly replace the equipment and either discard the old unit or refurbish it for eventual return to the pool. The Navy took the same approach toward some submarine equipment. After a few years the Navy assigned a reliability engineer to conduct the "age-reliability" procedures that airlines use to revise their equipment replacement schedules in light of experience.

Age-reliability analysis takes advantage of the fact that the schedules cannot be followed exactly and equipment sometimes fails early or stays in place beyond its planned replacement time. The analyst draws reliability curves for each type of equipment. The length and shape of the curve tell whether the equipment is being replaced too early or too late. Either case represents avoidable costs.

The reliability engineer was able to recommend revisions that offered significant savings with high statistical certainty of improving system readiness and safety. That was the easy part of the task. The difficult part was preparing the necessary data. He compared refurbishment facility records, equipment issue and inventory reports, and ship maintenance records, and found that often they did not match, requiring careful study and conference with submarine engineers and maintenance supervisors to reconstruct the equipment events, operating hours, and cost of repair.

Checking with the airlines, he discovered that they find fewer such errors because when the replacement schedules are instituted they include a data system to be used for subsequent analyses. Data collection discipline is high because it is associated with cost centers.

Lesson learned:

When preparing a preventive maintenance plan, include a data system to support later analysis of plan effectiveness and opportunities to improve system availability while significantly reducing costs.

3.5 Collecting Test Data

The specifications for a new avionics system emphasized reliability and maintainability. The client established a scoring board to review flight test data, perform reliability development and growth analyses, and recommend design improvements. The board members participated in planning the tests and preparing R&M data collection forms and instructions.

The early flight tests were conducted by the manufacturer, using system engineers to operate the equipment, perform maintenance, and collect data. But the board members found they couldn't score the results. The records did not consistently report the system operating hours, failure modes, and other information they needed. On the scoring board's recommendation, the client provided observers to collect data from all subsequent flight tests and engineering investigations. The results weren't perfect but the board could do its job of overseeing the growth of system reliability.

Lesson learned:

Don't depend on the manufacturer to collect test data. Have your own trained observers on site to make sure you get the data you'll need to do your job.

3.6 Data Byproducts and Fringe Benefits

When a reliability engineer was performing R&M predictions and allocations for a new system to be developed for a government agency, he noted that much of the same statistics had been researched from equipment histories before, in previous tasks for the same agency. New designs often include components from existing systems. He worked with the agency to create an R&M data bank that stored experience data from existing systems, for use not only in improving the current systems but also in designing reliability into future systems.

Lesson learned:

Valid R&M data contributes not only to the fielded product but also to future products. Treasure it.

Topic 4, Testing

*The only certainty about testing is Murphy's Law.
If something can go wrong, it will.*

All performance tests are of interest to the reliability engineer, not just the R&M-specific tests. They all accumulate operating hours and incur failures, and can contribute to reliability growth analyses and R&M reviews.

4.1 Reliability Test Reliability

A new-technology airborne system was being developed to detect enemy anti-aircraft weapons. The military services created a special project office to oversee testing against system specifications, including R&M specifications. Project office reliability engineers participated in the test planning.

In an early test the objective was to see if the system could reliably detect threats coming from all angles. R&M engineers and other experienced testers laid thorough plans for system preparation and maintenance, test procedures and data collection. Another agency had responsibility for simulating threat emissions.

The test results had limited usefulness because some of the emitters failed and there were no spares on hand.

Lesson learned:

R&M test planning is the business of preparing for every eventuality. One point that is easy to overlook is that when other organizations are involved they may not understand this principle. Review everybody's their plans to make sure they have also prepared for every eventuality.

4.2 Reliability Test Reliability, Revisited

The story in lesson 4.1 has a happy ending. In a subsequent test the system was installed in aircraft that flew past simulated threats. The objective was to see whether the system could reliably detect, pinpoint and classify them.

This time the R&M engineers and other planners had learned their lesson. They reviewed the plans made by the agency that provided the emitters, to make sure that testers and spares were on hand. Sure enough, one failed and was promptly replaced and checked for proper operation without interrupting the test.

Lesson learned:

See lesson 4.1.

4.3 Maintainability Demonstrators

An Army agency tasked a reliability engineer to observe and analyze the results of a maintainability demonstration to be conducted on a prototype avionic system. The engineer did not have high hopes for the usefulness of the demo. Based on previous experience, he expected that the agency conducting the demo might not be able to provide the necessary trained personnel, usually leading to a last-minute decision to perform the demo using factory engineers.

The Army Aviation Training Command performed the demonstration. The ATC selected and trained several soldiers from the appropriate specialty, including an anthropometric distribution of large, average-size and small male and female technicians. The soldiers performed system maintenance actions while wearing utility uniforms, flight suits and cold weather clothing.

In addition to making the necessary calculations for comparison against system maintainability specifications, the observer was able to recommend design improvements. For example in addition to other problems, the smallest maintainer could not release a squeeze-type clamp and the largest maintainer could not reach into a narrow access opening.

Lesson learned:

In a maintainability demonstration, it is important to convince the client to furnish adequate personnel and training resources for conducting the demo. Otherwise there is risk of buying maintenance headaches that could have been prevented before accepting the design.

4.4 Test Automation

To improve efficiency, a company's test lab automated their reliability test process, resulting in additional assets and fewer people. This had an adverse impact on earnings, because the contract paid for people, not equipment. Management instructed the test lab to get rid of the expensive automation equipment and hire more people. Not only did profits improve, but test quality also improved, and more people meant more ideas for improving the product.

Lesson learned:

Before deciding to automate a test process, carefully consider the impact on the quality of work.

4.5 Assumed Test Conditions

A company ran a 10-day accelerated life test on their product in an environmental chamber, using a wet-bulb/dry-bulb controller to maintain and record humidity. The wet bulb and dry bulb temperatures were almost exactly equal, and water droplets could be seen on the window, so they thought the humidity was at the required high percentage.

It turned out the chamber was as dry as a bone. The wet bulb was actually dry and the water drops were between the panes of glass in the window. They had to revise their procedures and re-run the 10-day test.

Lesson learned:

Examine test plans and procedures carefully. The hidden assumptions can ruin everything.

4.6 Test Figure of Merit

An engineer was assigned to take over the submission of availability reports on an emergency communications system. The specified requirement was “availability no less than 0.95” Previous reports had shown availability 0.99+ every month, but the system was rumored to be undependable. The engineer found that the monthly availability figures came from the vendor, who measured system downtime from the time the system was discovered to be down until vendor engineers reported it restored:

$$\text{Availability} = 1 - \frac{\text{Reported downtime}}{\text{Calendar time}} .$$

He began testing at irregular intervals, and calculated system availability as:

$$\text{Availability} = \frac{\text{Total successful trial time}}{\text{Total attempted trial time}} .$$

The resulting monthly availability averaged 0.65. Vendor payment was suspended until the system met R&M specifications.

Lesson learned:

In writing system R&M specifications, select figures of merit that relate directly to user requirements, and define how the figures are to be tested and calculated. Readiness specifications for systems that are seldom used may require special planning.

4.7 Stuff Happens, Sometimes for the Better

Flight testing of prototype equipment was inadequate to establish a starting point for reliability growth curves. The reliability engineer knew he was unlikely to be granted additional hours but prepared for it anyway. The next performance test required that two equipped aircraft fly cross-country to the test site and back. He noted that the trip amounted to more air time than all flight tests to date. He arranged to operate the equipment en route and accumulate the needed hours. In addition, an unexpected failure mode occurred that proved to be important.

Lesson learned:

In reliability testing, as in life, good luck springs from good planning.

Topic 5, Analysis

The worth of a system analysis depends, more than any other factor, on how well the analyst understands the system.

– Martin Binkin

5.1 Scope of the Analysis

A large system was installed at several remote sites. Each major subsystem was supported from a different central depot where parts were bought, stored, and repaired. The support manager tasked a reliability engineer to learn what computer models were used to optimize depot operations and on-site stocks, and recommend changes to improve overall system availability.

The engineer was familiar with the adage from Operations Research methodology that the first step should be to widen the scope of the analysis. He included the movement of parts between the depots and the sites, and also examined the standard planning factors the managers were using in their computer models.

The engineer reported that (a) The depot managers and on-site stock managers used several different computer models and all were effective in optimizing their local operations, (b) Some planning factors had been established long ago as “official” but were years out of date, and (c) Distribution of replacement parts from the depots and the

return of removed parts from the sites was unsystematic. Logistic support managers withheld and batched shipments in a way that was efficient from their viewpoint but degraded overall system availability.

Lesson learned:

In any system analysis, an initial step is to widen your view of what you were asked to do. If the client knew where the real problem was, he wouldn't have needed you.

5.2 Keeping Copious Notes

A reliability consultant contracted to perform an R&M prediction for a new-technology medical system. He began keeping a detailed log of people he talked to, references he used and decisions he made. He found the log made the work go more smoothly, prevented some mistakes, helped him answer questions professionally, and assisted documenting the report. A year later he was asked for R&M predictions on planned new versions of the system. The log paid off again by guiding revision of the system configuration and updating the parts R&M data.

Lesson learned:

Take a minute to write it down. You'll be damned glad you did.

5.3 Spread Sheet Mission Simulation

The Navy was considering replacing the engines in a class of ships. A reliability engineer was asked to prepare mission operational availability curves, to be used in establishing R&M specifications for replacement engines and in supporting trade studies. It was initially assumed he would use the Navy's Monte Carlo (open form) mission simulation software. However, smooth curves would obviously be required and he knew that the term "smooth Monte Carlo curve" is an oxymoron. But the situation was well suited for closed-form simulation, and didn't need special mission simulation software.

He created a spread sheet that directly displayed the figure of merit, in this case families of mission availability curves against equipment R&M requirements, without the characteristic Monte Carlo wobble. The smooth curves made it easy for the client to plot the engineering options and make cost/performance comparisons.

In another case a marine engineering firm hired a consultant to predict mission operational availability of a new-technology propulsion system. The model was to be used to compare design alternatives and support equipment selection. The company planned to run it in a company-owned Monte Carlo simulator. The consultant noted that although the system configuration was extensive, it was well suited for closed form

simulation. He developed a spread sheet model that was easier to use and better suited for the purpose.

Lesson learned:

In constructing an R&M model that may be used for comparing alternatives, keep an open mind to closed-form simulation. Closed-form equations do not have the flexibility of the Monte Carlo open form, but neither do they introduce a random element.

5.4 Warranties

A Navy agency assigned a reliability engineer to study the use and cost effectiveness of equipment warranties in commercial ships.

The engineer interviewed commercial shippers and marine equipment manufacturers. He learned that (a) Warranties for Navy equipment would probably expire because the Navy makes advance purchases and stocks ready spares for extended periods. And (b) Warranties are seldom enforced because they are given to the ship Captain, who has no incentive to enforce them.

The engineer reported on methods being used to calculate warranty cost effectiveness, and advised that warranties might apply to vendor shipments for immediate use, but then only if the ship's officers were given incentive to enforce them.

Lesson learned:

Reliability warranties cannot be cost effective until procedures are established to ensure their conditions are met, and the official who holds the warranties is given incentive to enforce them.

5.4 Analysis Figures of Merit

At a symposium a professor presented a paper on risk analysis. The example was the shipment of frozen food containers that were electrically powered by the shipper. The analysis was based on the risk that power would fail.

The first comment from the floor was that the figure of merit was not appropriate. The mission of the shipper's electric plant was not to provide uninterrupted power, it was to maintain container temperatures below a specified limit. It would be more appropriate to first determine the length of time a container could tolerate the loss of power and continue to meet temperature specifications. The figure of merit would be the risk that a power failure exceeded the allowed time.

In another case a military support facility wanted to improve their performance. They hired a consulting firm to help them focus on critical line items. The firm initiated a program of measuring overall performance as the average availability of all supported systems. They periodically calculated availability figures and identified those systems with the worst performance, for action.

Reliability engineers felt the emphasis was being misplaced on optimizing system availability from the facility's viewpoint, instead of the facility's real mission, which was to support military readiness. The measure of facility performance should be weighted by each system's contribution to the readiness of the military units being supported.

Lesson learned:

Investigate thoroughly with the client before establishing the objective and figures of merit for analysis. The risk is spending time and resources to obtain results that do not satisfy either of you.

Topic 6, Analysis Results

Experience tells us that an R&M analysis report is unlikely get action for system improvement if the analyst runs out of time, money and ideas all at the same time.

6.1 Engineering Recommendations

In a recent brochure advertising R&M analytical software, the vendor said, "Once his data has been processed by the [FMECA] program, the engineer can kick back and relax. His job is complete." This statement was obviously not written by a reliability engineer. In fact, cases where system R&M improvement was obtained by presenting raw analysis results are very rare. The reliability engineer stands a far better chance of getting action if he offers specific engineering recommendations.

To show how this works, in planning a reliability improvement study for a large organization, the reliability analysts included a final step in their task schedule, to prepare recommendations. During most of the analysis, they received only limited cooperation from the corporate engineers, who were not accustomed to considering quantitative R&M analysis in their planning. The reliability analysts prepared a draft report in which they translated their results into a list of specific engineering recommendations for system redesign, equipment selection, special studies, etc.

The draft report got the organization's engineers' full attention because it shifted the focus from whether they needed to read the report to how they should respond to it. They pointed out:

- Additional data the analysts had not been given,
- Additional engineering alternatives to overcome some R&M shortfalls, and
- Additional justification for action.

With this help, the analysts completed a report recommending 14 R&M actions. Seven of them were approved, which may be some kind of record.

Lesson learned:

When scheduling an analysis task, include a final step to draft specific engineering recommendations, review them with client engineers, and finalize them.

You can download a PDF file of Reid Willis' lessons learned by [clicking here](#).