

Blueprint for a Comprehensive Reliability Program

ReliaSoft Corporation R & D Reports

Author:	Smith, Crawford
Last Revised on:	February 3, 2003

Prepared by:

ReliaSoft Corporation

Phone: +1.520.886.0410

Fax: +1.520.886.0399

www.Reliasoft.com

TABLE OF CONTENTS

1	Introduction	1
2	Foundations of Reliability	2
2.1	Fostering a Culture of Reliability	2
2.2	Product Mission	3
2.2.1	Reliability Specifications	3
2.3	Universal Failure Definitions	4
3	Reliability Testing	6
3.1	Customer Usage Profiling	6
3.2	Test Types	6
3.2.1	Development Testing	7
3.2.2	Manufacturing Testing	8
4	Field Data.....	10
4.1	The “Disconnect” Between In-House and Field Data.....	10
4.2	Sales and Forecasting (Shipping) Data	10
4.3	Warranty Data.....	11
4.4	Field Service Data.....	11
4.5	Customer Support Data	12
4.6	Returned Parts/Failure Analysis Data.....	12
5	Data Collection.....	13
5.1	In-House Test Data Collection	13
5.1.1	Test Log.....	14
5.1.2	Failure Log.....	14
5.1.3	Service Log.....	15
5.2	Field Data Collection.....	16
5.2.1	ReliaSoft’s Dashboard	16
6	Data Analysis and Reporting	19
6.1	Non-Parametric Analysis	19
6.1.1	Non-Parametric Reliability Analysis.....	19
6.2	Parametric Analysis	21
6.2.1	Examples of Reporting for Parametric Data Analysis	22
6.2.1.1	Probability Plot.....	22
6.2.1.2	Reliability Function	23
6.2.1.3	Probability Density Function	23
6.2.1.4	Failure Rate Function	23
6.2.1.5	Life vs. Stress	24
6.2.1.6	Likelihood Function.....	24
6.2.1.7	Reliability Importance	25
6.2.1.8	Reliability Growth.....	25
7	Bringing It All Together	26
7.1	Connecting Field and Lab Data	26
7.2	Reliability Growth	26
7.3	Optimum Design Level Determination	27
7.4	Competitive Assessment	28
7.5	Marketing and Advertising	28

1 INTRODUCTION

At the highest level, the purpose of a reliability engineering program is to test and report on the reliability of an organization's products. This information is then used to assess the financial impact of the reliability of the products, and to improve the overall product reliability and consequently the financial strength of the organization.

Reliability assessment is based on the results of testing from in-house labs and data pertaining to the performance results of the product in the field. The data produced by these sources is to be utilized to accurately measure and improve the reliability of the products being produced. This is particularly important as market concerns drive a constant push for cost reduction. However, one must be able to keep a perspective on "the big picture," instead of merely looking for the quick fix. It is often the temptation to cut corners and save initial costs by using cheaper parts or cutting testing programs. Unfortunately, cheaper parts are usually less reliable, and inadequate testing programs can allow products with undiscovered flaws to get out into the field. A quick savings in the short term by the use of cheaper components or small test sample sizes will usually result in higher long-term costs in the form of warranty costs, or loss of customer confidence. The proper balance must be struck between reliability, customer satisfaction, time to market, sales and features. Figure 1 illustrates this concept. The polygon on the left represents a properly balanced project. The polygon on the right represents a project in which reliability and customer satisfaction have been sacrificed for the sake of sales and time to market.




Figure 1 - Graphical Representation of Balanced and Unbalanced Projects.

As the figure above illustrates, a proper balance must be struck between product reliability and the other aspects of the business, including time-to-market, manufacturing costs, sales, product features, and customer satisfaction. Through proper testing and analysis in the in-house testing labs, as well as collection of adequate and meaningful data on a product's performance in the field, the reliability of any product can be measured, tracked and improved, leading to a balanced organization with a financially healthy outlook for the future.

2 FOUNDATIONS OF RELIABILITY

There are a certain number of "up-front" activities that need to be undertaken in order to be able to successfully implement a reliability program. Most of these activities are relatively simple, but they are vital to the design and implementation of a reliability program. It is like the foundation of a house: relatively inexpensive and largely forgotten once the major construction has begun, but vital to the structure nonetheless. The reliability "foundation building" concepts need to be addressed in order to put a strong reliability program in place.

2.1 FOSTERING A CULTURE OF RELIABILITY

The most important part of developing a reliability program is having a culture of reliability in the organization. It is vital that everyone involved in the product's production, from upper management to assembly personnel, understands that a sound reliability program is necessary for the organization's success.

Achieving this culture of reliability may actually be more difficult than it seems, as some organizations may not have the history or background that lends itself to the support of a reliability program. This can be particularly true in situations where the organization has had a niche market or little or no previous competition with the products that it produces. In the past, the organization's customers may have had to accept the reliability of the product, good or bad. As a consequence, the organization may have developed a mentality that tends to overlook the reliability of a product in favor of the "damn-the-torpedoes, full-steam-ahead" method of product development. In this type of organization, reliability engineering methods and practices tend to be viewed as superfluous or even wasteful. "We don't need all of this reliability stuff, we'll just find the problems and fix them," tends to be the attitude in these circumstances. Unfortunately, this attitude often results in poorly-tested, unreliable products being shipped to customers.

The first step in developing the necessary culture of reliability is to have the support of the organization's top management. Without this, the implementation of a reliability program will be a difficult and frustrating process. Adequate education as to the benefits of a properly constructed reliability program will go a long way towards building support in upper management for reliability techniques and processes. Most important is the emphasis of the financial benefits that will accrue from a good reliability program, particularly in the form of decreased warranty costs and increased customer goodwill. This latter aspect of the benefits of reliability engineering can sometimes be an elusive concept to appreciate. An adage in the reliability field states that if customers are pleased with a product, they will tell eight other people, but if they are dissatisfied, they will tell twenty-two other people.¹ While this anecdote is rather eye-opening, it must be put in a financial context to have the full impact. It is possible to construct a model that links reliability levels of a product with the probability of repeat sales. It is therefore possible to calculate a loss of sales revenue based on the unreliability of the product. This type of information is useful in educating upper management on the financial importance of reliability. Once the upper management has been adequately educated and is supportive of the

¹ This adage is rather dated. Given the popularity of the Internet and the proliferation of newsgroups and web sites dedicated to disgruntled customers and employees, this number is probably even higher than twenty-two.

implementation of reliability concepts, it will be a great deal easier to go about implementing those concepts.

However, one should not stop with upper management when it comes to educating an organization on the benefits of a proposed reliability program. It is vital to have the support and understanding of the rest of the organization as well. Since the implementation of a reliability program will affect the day-to-day activities of middle management, the engineers, and the technicians, it is also necessary to convince these groups of the benefits of reliability-oriented activities. It is important to demonstrate to them how these activities, which may initially seem pointless or counterproductive to them, will ultimately benefit the organization. For example, if test technicians have a good understanding of how the test data are going to be put to use, they will be less likely to cut corners while performing the test and recording the data. Overall, a reliability program stands the greatest chance of success if everyone in the organization understands the benefits and supports the techniques involved.

2.2 PRODUCT MISSION

The underlying concept in characterizing the reliability of a product involves the concept of product mission (e.g., operate for 36 months, or complete 1000 cycles). A textbook definition of reliability is:

*"The conditional probability, at a given confidence level, that the equipment will perform its intended functions satisfactorily or without failure, i.e., within specified performance limits, at a given age, for a specified length of time, function period, or mission time, when used in the manner and for the purpose intended while operating under the specified application and operation environments with their associated stress levels."*²

With all of the conditions removed, this boils down to defining reliability as the ability of a product to perform its intended mission without failing. The definition of reliability springs directly from the product mission, in that product failure is the inability of the product to perform its defined mission.

2.2.1 RELIABILITY SPECIFICATIONS

In order to develop a good reliability program for a product, the product must have good reliability specifications. These specifications should address most, if not all, of the conditions in the reliability definition above, including mission time, usage limitations, operating environment, etc. In many instances, this will require a detailed description of how the product is expected to perform reliability-wise. Use of a single metric, such as MTBF, as the sole reliability metric is inadequate. Even worse is the specification that a product will be "no worse" than the previous model. An ambiguous reliability specification leaves a great deal of room for error, and this can result in a poorly-understood and unreliable product reaching the field.

Of course, there may be situations in which an organization lacks the reliability background or history to easily define specifications for a product's reliability. In these instances, an analysis of existing data from previous products may be necessary. If enough information exists to characterize the reliability performance of a previous product, it should be a relatively simple matter to transform this

² Kececioglu, Dimitri, *Reliability Engineering Handbook*, Vol. 1, Prentice-Hall, 1991.

historical product reliability characterization into specifications of the desired reliability performance of the new product.

Financial concerns will definitely have to be taken into account when formulating reliability specifications. Planning for warranty and production part costs is a significant part of financial planning for the release of a new product. Based on financial inputs such as these, a picture of the required reliability for a new product can be established. However, financial wishful thinking should not be the sole determinant of the reliability specifications. It can lead to problems such as unrealistic goals, specifications that change on a regular basis to fit test results, or test results that get "fudged" in order to conform to unrealistic expectations. It is necessary to couple the financial goals of the product with a good understanding of product performance in order to get a realistic specification for product reliability. A proper balance of financial goals and realistic performance expectations are necessary to develop a detailed and balanced reliability specification.

2.3 UNIVERSAL FAILURE DEFINITIONS

Another important foundation for a reliability program is the development of universally agreed-upon definitions of product failure. This may seem a bit silly, in that it should be fairly obvious whether a product has failed or not,³ but such a definition is quite necessary for a number of different reasons.

One of the most important reasons is that different groups within the organization may have different definitions as to what sort of behavior actually constitutes a failure. This is often the case when comparing the different practices of design and manufacturing engineering groups. Identical tests performed on the same product by these groups may produce radically different results simply because the two groups have different definitions of product failure. In order for a reliability program to be effective, there must be a commonly accepted definition of failure for the entire organization. Of course, this definition may require a little flexibility depending on the type of product, development phase, etc., but as long as everyone is familiar with the commonly accepted definition of failure, communications will be more effective and the reliability program will be easier to manage.

Another benefit of having universally agreed-upon failure definitions is that it will minimize the tendency to rationalize away failures on certain tests. This can be a problem, particularly during product development, as engineers and managers may tend to overlook or diminish the importance of failure modes that are unfamiliar or not easily replicable. This tendency is only human, and a person who has spent a great deal of time developing a product may find justification for writing off oddball failures as a "glitch" or as failure due to some other external error. However, this type of mentality also results in products with poorly-defined but very real failure modes being released into the field. Having a specific failure definition that applies to all or most types of tests will help to alleviate this problem.

However, a degree of flexibility is called for in the definition of failure, particularly with complex products that may have a number of distinct failure modes. For this

³ This is closely related to the concept of product mission. A good baseline definition of failure is the inability of the product to perform its mission. From that basic definition, more detailed categorizations of failure can be developed.

reason, it may be advisable to have a multi-tiered failure definition structure that can accommodate the behavioral vagaries of complex equipment.

The following three-level list of failure categories is provided as an example:

- Type I – *Failure*: Severe operational incidents that would definitely result in a service call, such as part failures, unrecoverable equipment hangs, DOAs, consumables that fail/deplete before their specified life, onset of noise, and other critical problems. These constitute “hard-core” failure modes that would require the services of a trained repair technician to recover.
- Type II – *Intervention*: Any unplanned occurrence or failure of product mission that requires the user to manually adjust or otherwise intervene with the product or its output. These tend to be “nuisance failures” that can be recovered by the customer, or with the aid of phone support. Depending on the nature of the failure mode, groups of the Type II failures could be upgraded to Type I if they exceed a predefined frequency of occurrence.
- Type III – *Event*: Events will include all other occurrences that do not fall into either of the categories above. This might include events that cannot directly be classified as failures, but whose frequency is of engineering interest and would be appropriate for statistical analysis. Examples include failures caused by test equipment malfunction or operator error.

During testing, all of these occurrences will be logged with codes to separate the three failure types. Other test-process-related issues such as deviations from test plans will be logged in a separate test log. There should be a timely review of logged occurrences to insure proper classification prior to metric calculation and reporting. These and other procedural issues will be discussed in detail in the next section.

3 RELIABILITY TESTING

Reliability testing is the cornerstone of a reliability engineering program. It provides the most detailed forms of data in that the conditions under which the data are collected can be carefully controlled and monitored. Furthermore, the reliability tests can be designed to uncover particular suspected failure modes and other problems. The type of reliability testing a product undergoes will change along different points of its life cycle, but the overriding goal is to insure that data from all or most of the tests were generated under similar enough conditions so that an "apples-to-apples" comparison can be made of the product's reliability characteristics at different points in the product's life. It is for this reason that consistent and thorough reliability specifications and a standard definition of failure are up-front requirements to implementing reliability testing.

A properly designed series of tests, particularly during the product's earlier design stages, can generate data that would be useful in the implementation of a reliability growth tracking program. This will provide information helpful in making management decisions regarding scheduling, development cost projections and so forth. This information will also be useful in planning the development cycle of future products.

3.1 CUSTOMER USAGE PROFILING

An important requirement for designing useful reliability tests is to have a good idea of how the product is actually going to be used in the field. The tests should be based on a realistic expectation of the customer usage, rather than estimates or "gut feelings" about the way the customer will use the product. Tests based on mere speculation may result in a product that has not been rigorously tested and consequently may run into operational difficulties due to use stress levels being higher than anticipated. On the other hand, tests that are designed with a strong basis of information on how the product will be used will be more realistic and result in an optimized design that will exhibit fewer failures in the field.

Customer usage profiles can be designed to actively gather information on how the customers are actually using an organization's product. This design can range from a simple questionnaire to a sophisticated instrumentation within the product that feeds back detailed information about its operation. An incentive is often useful to get customers to sign on for a usage measurement program, particularly if it is an intrusive process that involves the installation of data collection equipment. Additionally, customers are often eager to participate in these programs in the knowledge that the information that they provide will ultimately result in a more reliable and user-friendly product.

3.2 TEST TYPES

In many cases, the type of testing that a product undergoes will change as the product's design becomes mature and the product moves from the initial design stages to final design release and production. Nevertheless, it is a good practice to continue to collect internally-generated data concerning the product's reliability performance throughout the life cycle of the product. This will strengthen the reliability growth analysis and help provide correlation between internal test results and field data. A brief summary of the various types of reliability tests is presented next.

3.2.1 DEVELOPMENT TESTING

Development testing occurs during the early phases of the product's life cycle, usually from project inception to product design release. It is vital to be able to characterize the reliability of the product as it progresses through its initial design stages so that the reliability specifications will be met by the time the product is ready for release. With a multitude of design stages and changes that could affect the product's reliability, it is necessary to closely monitor how the product's reliability grows and changes as the product design matures. There are a number of different test types that can be run during this phase of a product's life cycle to provide useful reliability information:

- **Component-level Testing:** Although component-level testing can continue throughout the development phase of a product, it is most likely to occur very early in the process. This may be due to the unavailability of parts in the early stages of the development program. There may also be special interest in the performance of a specific component if it has been radically redesigned, or if there is a separate or individual reliability specification for that component. In many cases, component-level testing is undertaken to begin characterizing a product's reliability even though full system-level test units are unavailable or prohibitively expensive. However, system-level reliability characterization can be achieved through component-level testing. This is possible if sufficient understanding exists to characterize the interaction of the components. If this is the case, the system-level reliability can be modeled based on the configuration of components and the result of component reliability testing.
- **System-level Testing:** Although the results of component-level tests can be used to characterize the reliability of the entire system, the ideal approach is to test the entire system, particularly if that is how the reliability is specified. That is, if the technical specifications call out a reliability goal for a specific system or configuration of components, that entire system or configuration should be tested to compare the actual performance with the stated goal. Although early system-level test units may be difficult to obtain, it is advisable to perform reliability tests at the system level as early in the development process as possible. At the very least, comprehensive system-level testing should be performed immediately prior to the product's release for manufacturing, in order to verify design reliability. During such system-level reliability testing, the units under test should be from a homogeneous population and should be devoted solely to the specific reliability test. The results of the reliability test could be skewed or confounded by "piggybacking" other tests along with it, and this practice should be avoided. A properly conducted system-level reliability test will be able to provide valuable engineering information above and beyond the raw reliability data.
- **Environmental and Accelerated Testing:** It may be necessary in some cases to institute a series of tests in which the system is tested at extreme environmental conditions, or with other stress factors accelerated above the normal levels of use. It may be that the product would not normally fail within the time constraints of the test, and, in order to get meaningful data within a reasonable time, the stress factors must be accelerated. In other cases, it may be

necessary to simulate different operating environments based on where the product will be sold or operated. Regardless of the cause, tests like these should be designed, implemented and analyzed with care. Depending on the nature of the accelerating stress factors, it is easy to draw incorrect conclusions from the results of these tests. A good understanding of the proper accelerating stresses and the design limits of the product are necessary to be able to implement a meaningful accelerated reliability test. For example, one would not want to design an accelerated test that would overstress the product and introduce failure modes that would not normally be encountered in the field. Given that there have been a lot of incredible claims about the capability of accelerated testing and the improbably high acceleration factors that can supposedly be produced, care needs to be taken when setting up this type of reliability testing program.

- **Shipping Tests:** Although shipping tests do not necessarily qualify as reliability tests *per se*, shipping tests or simulations designed to test the impact on the product of shipping and handling should be a prerequisite to reliability testing. This is because the effects of shipping will often have an impact on the reliability of the product as experienced by the customer. As such, it may be useful to incorporate shipping tests alongside the normal reliability testing. For example, it may be a good idea to put the units of a final design release reliability test through a non-destructive shipping test prior to the actual reliability testing in order to better simulate actual use conditions.

3.2.2 MANUFACTURING TESTING

The testing that takes place after a product design has been released for production generally tends to measure the manufacturing process rather than the product, under the assumption that the released product design is final and good. However, this is not necessarily the case, as post-release design changes or feature additions are not uncommon. It is still possible to obtain useful reliability information from manufacturing testing without diluting any of the process-oriented information that these tests are designed to produce.

- **Functionality Testing and Burn-In:** This type of testing usually falls under the category of operation verification. In these tests, a large proportion, if not all, of the products coming off of the assembly line are put on a very short test in order to verify that they are functioning. In some situations, they may be run for a predetermined "burn-in" time in order to weed out those units that would have early infantile failures in the field. Although it may not be possible to collect detailed reliability information from this type of testing, what is lost in quality is made up for in quantity. With the proper structuring, these tests can provide a fairly good picture of early-life reliability behavior of the product.
- **Extended Post-Production Testing:** This type of testing usually gets implemented near the end or shortly after the product design is released to production. It is useful to structure these types of tests to be identical to the final reliability verification tests conducted at the end of the design phase. The purpose of these tests is to assess the effects of the production process on the reliability of the product. In many cases, the test units that undergo reliability testing prior to the onset of actual production are hand-built or carefully adjusted prior to

the beginning of the reliability tests. By replicating these tests with actual production units, potential problems in the manufacturing process can be identified before many units are shipped.

- **Design/Process Change Verification:** This type of testing is similar to the extended post-production testing in that it should closely emulate the reliability verification testing that takes place at the end of the design phase. This type of testing should occur at regular intervals during production, or immediately following a post-release design change or a change in the manufacturing process. These changes can have a potentially large effect on the reliability of the product, and these tests should be adequate, in terms of duration and sample size, to detect such changes.

4 FIELD DATA

While reliability testing is vital to the implementation of a reliability program, it is not the sole source of product reliability performance data. Indeed, the information received from the field is the "true" measure of product performance, and is directly linked to the financial aspects of a product. In fact, a significant proportion of field data may be more finance-related than reliability-related. However, given the importance of the link between reliability and income, it is important to insure that adequate reliability information can be gleaned from field performance data. In many cases, it is not too difficult to adapt field data collection programs to include information that is directly applicable to reliability reporting.

Some of the most prevalent types of field data are: sales and forecasting data, warranty data, field service data, customer support data, and returned parts/failure analysis data, as discussed below. These discussions will tend towards generalizations, as every organization has different methods of monitoring the performance of its products once they are in the field. However, the illustrations below give a good general overview of how different types of field data may be collected and put to use for a reliability program.

4.1 THE "DISCONNECT" BETWEEN IN-HOUSE AND FIELD DATA

It should be noted at this point that there will usually be a "disconnect," or seeming lack of correlation, between the reliability performance of the products in the field and the results of in-house reliability testing. A typical rule of thumb is to expect the reliability in the field to be half of what was observed in the lab. Some of the specific causes of this disparity are discussed below, but in general the product will usually receive harsher treatment in the field than in the lab. Units being tested in the labs are often hand-built or carefully set up and adjusted by engineers prior to the beginning of the test. Furthermore, the tests are performed by trained technicians who are adept at operating the product being tested. Most end-use customers do not have the advantage of a fine-tuned unit and training and experience in its operation, thus leading to many more operator-induced failures than were experienced during in-house testing. Also, final production units are subject to manufacturing variation and transportation damage that test units might not undergo, leading to yet more field failures that would not be experienced in the lab. Finally, the nature of the data that goes into the calculations will be different; in-house reliability data is usually a great deal more detailed than the catch-as-catch-can data that characterizes a great deal of field data.⁴ As can be seen, there are any number of sources for the variation between field reliability data and in-house reliability test results. However, with careful monitoring and analysis of both sources of data, it should be possible to model the relationship between the two, allowing for more accurate prediction of field performance based on reliability testing results.

4.2 SALES AND FORECASTING (SHIPPING) DATA

The sales and forecasting category of data is a sort of general-use data that is necessary as a basis for many other analyses of field data. Essentially, this information provides the analyst with a figure for the population of products in the

⁴ This problem can be ameliorated by setting up a comprehensive field data collection system. This is discussed in more detail in the section on Field Data Collection.

field. Knowing how many units are being used at any given time period is absolutely vital to performing any sort of reliability-oriented calculations. Having an accurate measurement of the number of failures in the field is basically useless if there is not a good figure for the total number of units in the field at that time.

4.3 WARRANTY DATA

The Warranty Data category is somewhat of a catch-all category that may or may not include the other types of field data listed below, and may not contain adequate information to track reliability-related data. Since most warranty systems are designed to track finances and not performance data, some types of warranty data may have very little use for reliability purposes. However, it may be possible to acquire adequate reliability information based on the inputs of the warranty data, if not the actual warranty data itself. For example, a warranty system may have ship dates and service call dates, but not actual time-to-failure data. In this case, we must make the assumption that the failure time is approximately equal to the difference between the ship date and service call date, even though the product may not have actually been used during the extent of that time before it failed. This of course is a case of "garbage in, garbage out," and a poorly designed warranty tracking system will yield poor or misleading data regarding the reliability of the product. At the very least, there should be a degree of confidence regarding the raw number of failures or warranty hits during a particular time period. This, coupled with accurate shipping data, will allow a crude approximation of reliability based on the number of units that failed versus the number of units operating in the field in any given time period.

4.4 FIELD SERVICE DATA

Field Service data is a category of data that is connected with field service calls where a repair technician manually repairs a failed product during an on-site visit. This is a potentially powerful source of field reliability information, if a system is in place to gather the necessary data during the service call. However, the job of the service technician is to restore the customer's equipment to operating condition as quickly as possible, and not necessarily to perform a detailed failure analysis. This can lead to a number of problems. First, the service technician may not be recording information necessary to reliability analysis, such as how much time the product accumulated before it failed. Second, the technician may take a "scattershot" approach to repair. That is, based on the failure symptom, the technician will replace all of the parts whose failure may result in the failure of that particular system. It may be that only one of the parts that were replaced had actually failed, so it is necessary to perform a failure analysis on all of the parts to determine which one was actually the cause of the product failure. Unfortunately, this is not always done, and if it is, the parts that have had no problem found with them will often be returned to field service circulation. This may lead to another potential source of error in field service data, in that used parts with unknown amounts of accumulated time and damage may be used as replacement parts on subsequent service calls. This makes tracking and characterizing field reliability very difficult. From a reliability perspective, it is always best to record necessary failure information, avoid using the "scattershot" approach to servicing failed equipment, and always use new units when making part replacements.

4.5 CUSTOMER SUPPORT DATA

Customer Support data is a category of data that comes from phone-in customer support services. In many cases, it may be directly related to the field service data in that the customer with a failed product will call to inform the organization. In some circumstances, it may be possible to solve the customer's problem over the phone, or adequately diagnose the cause of the problem so that a replacement part may be sent directly to the customer without requiring a service technician to make an on-site visit. Ideally, the customer support and field service data would reside in the same database, but this is not always the case. Regardless of the location, customer support data must always be screened with care, as the information does not always reflect actual problems with the product. Many customer support calls may concern usability issues or other instances of the customer not being able to properly use the product. In cases such as this, there will be a cost to the organization or warranty hit, even though there is no real fault or failure for the product. For example, a product that is very reliable, but has a poorly written user manual may generate a great deal of customer support calls. This is because, even though the product is working perfectly, the customers are having difficulty operating the product. This is a good example of one of the sources of the "disconnect" between in-house and field reliability data.

4.6 RETURNED PARTS/FAILURE ANALYSIS DATA

As was mentioned earlier, failed parts or systems are sometimes returned for more detailed failure analysis than can be provided by the field service technician. Data from this area are usually more detailed regarding the cause of failure, and are usually more useful to design or process engineers than to reliability engineers. However, it is still an important source of information regarding the reliability behavior of the product. This is especially true if the field service technicians are using the "scattershot" approach to servicing the failed product, replacing a number of parts that may or may not be defective. If this is the case, it is necessary for all of the returned parts to be analyzed to determine the true cause of the failure. The results of the failure analysis should be correlated with the field service records in order to provide a complete picture of the nature of the failure. Often this correlation does not occur, or the returned parts are not analyzed in a timely fashion. Even if the analysis is performed correctly, there tend to be a significant proportion of returned parts with which no problem can be found. This is another example of a potential cause of the disparity between lab and field reliability data. However, even if the failure analysis group is unable to assign a cause to the failure, a failure has taken place, and the organization has taken a warranty hit. In the field, the performance the customer experiences is the final arbiter of the reliability of the product.

5 DATA COLLECTION

Data collection is the framework for a good reliability engineering program. It is necessary to have an accurate and comprehensive system of recording data relating to a product's reliability performance in order to be able to produce meaningful reliability reports. Although the nature of the data collection system will differ based on the type of data being collected, there must be a certain number of common elements in the type of data being collected and the way the information is recorded. This is necessary in order to provide the continuity necessary for developing an accurate assessment of the "cradle-to-grave" reliability of a product.

Whenever possible, computers should be employed in the data collection and recording process. Of course, the method of data collection will vary with the product under consideration, but given the decreasing cost and increasing power of computer systems, it should not be very difficult to set up a computerized data collection system. In some cases, it is even possible to automate much of the data collection process, thus further decreasing the potential for data recording errors. The concepts of in-house data collection, field data collection, and ReliaSoft's Dashboard are presented in more detail below.

5.1 IN-HOUSE TEST DATA COLLECTION

In a previous section, the different types of in-house reliability testing were discussed. One of the most important aspects of setting up an in-house reliability testing program lies in having certain common elements that extend across all of the different types of tests. This can be aided greatly by a degree of uniformity in the data collection process. By having a core group of data types that are collected from every test that takes place, it is easier to perform similar analyses across a variety of different test types. This lends a great deal of continuity in test analysis and reporting that will benefit the reliability program and the entire organization.

As mentioned earlier, it is highly beneficial to automate the data collection and recording process wherever possible. The method of data collection will differ from product to product, depending on whether it is possible to use a computer interface to operate the product during the test and automatically record the test results. Regardless of the method employed in running the test, it is always possible and advisable to use a database system to keep track of the test results. The use of relational databases makes it fairly easy to manipulate large quantities of data, and greatly aids in the reporting process. Properly managed, it is even possible to automate some if not all of the data reporting process using database information. Of course, human oversight is always a necessity when dealing with data analysis of any type, but proper use of database structuring and manipulation can make the reliability engineer's job much easier.

In setting up a database system for collecting data from in-house reliability testing, there are a minimum number of data types that need to be included in the database structure. For the purposes of in-house reliability data collection, it is recommended to have at least three related databases: a test log, a failure log, and a service log. Detailed descriptions of these databases and the information they should contain appear below.

5.1.1 TEST LOG

The test log contains detailed information on the tests being run on the products. The structure of the database will vary depending on the testing procedures and depending on the type of products for which data are being captured. If the product requires a test in which the test units are essentially just turned on and left to run until they fail or the time of the test expires, the test log will be fairly simple. However, if the product requires a variety of different inputs in order to be properly exercised during testing, the test log should be detailed enough to record all of the pertinent information. A suggested list of fields for the test log include:

- **Transaction number:** a unique identification code for the test log entry.
- **Test start date:** the date the test starts.
- **Test start time:** the time the test starts.
- **Test name:** the name or identifier for the test being run.
- **Test stage or step:** if the test is run in a series of stages or steps with different inputs, this field should provide a description or count of which segment of the test is being run, e.g. "Step 2," "High Temperature," etc.
- **Test inputs:** this describes the test inputs at each stage or step of the test. Depending on the nature of the product and the testing, it may be necessary to create a separate log for the specific steps and inputs of the test in order to keep the test log from being too cluttered with specific step/input information.
- **Operator comments:** specific comments regarding the test that may be useful when performing subsequent analyses.
- **Deviations:** descriptions of deviations from the original test plan.

5.1.2 FAILURE LOG

The failure log is where the majority of the information that is important to the generation of reliability results will reside. Care should be taken in the construction of this database so that all of the pertinent failure information will be collected every time a failure occurs. At the same time, it should not have so many fields as to be unwieldy when conducting a reliability analysis. This might slow down the overall testing process if a large amount of minutely detailed information needs to be recorded. Developing a method of automating the data collection process will alleviate this problem, but that is not always possible. A suggested list of fields for the failure log include:

- **Transaction number:** a unique identification code for the failure log entry.
- **Test log cross-reference:** the transaction number for the test log entry which corresponds to the test on which the failure occurred.
- **Service log cross-reference:** the transaction number for the most recent service log entry.
- **Failure date:** the date when the failure occurred.
- **Failure time:** the time when the failure occurred.

- **Failure type:** this describes the failure type encountered, particularly if a multi-tiered system of failure classification is being used.
- **Test stage:** the stage or step of the test when the failure occurred. This can be cross-referenced to the appropriate test or input log.
- **Symptom code:** this is the symptom noticed by the operator when the failure occurred. The type of symptom code can be cross-linked to a preliminary failure code.
- **Failure code:** this is a code that describes the actual mode of failure. A preliminary failure code can be generated based on the symptom code, but a failure analysis engineer should make the final disposition of the failure code.
- **Failed part ID:** this describes the part or parts that caused the failure. If possible, the failed part serial number should be included as a separate field.
- **Resolution:** this field describes what action was taken to restore the failed unit to operational status. This field should be cross-linked to the service log.
- **Comments:** specific comments regarding the failure that may be useful when performing subsequent analyses.

5.1.3 SERVICE LOG

The purpose of the service log is to track and record any service actions or modifications performed on test units. It is important to keep a concise record of any service actions performed on test units because even a relatively small modification or repair can potentially have a large effect on the performance of the test units. By requiring service technicians and engineers to use a service log whenever they work on a test unit, the amount of unofficial "tinkering" with a system will be minimized, thus reducing unexplained changes in test unit performance. A service log entry should be made whenever a test unit is installed or upgraded. This allows for tracking design level or version number changes across the tests. A suggested list of fields for the service log include:

- **Transaction number:** a unique identification code for the service log entry.
- **Test log cross-reference:** the transaction number for the test log entry which corresponds to the test during which the service was performed.
- **Service date:** the date on which the service was performed.
- **Service time:** the time at which the service was performed.
- **Current version identifier:** identifies the revision or design level of the test unit before the service is performed.
- **New version identifier:** identifies the revision or design level of the test unit after the service is performed. This will be the same as the current version identifier unless the service performed upgrades the test unit to the next level.
- **Service type:** describes the service performed.
- **Part modified/replaced:** a description/serial number of the part modified or replaced during the service action.

- **Comments:** specific comments regarding the service action that may be useful when performing subsequent analyses.

5.2 FIELD DATA COLLECTION

Depending on the circumstances, collection of field data for reliability analyses can be either a simple matter or major headache. Even if there is not a formal field data collection system in place, odds are that much of the necessary general information is being collected already in order to track warranty costs, financial information, etc.. The potential drawback is that the data collection system may not be set up to collect all of the types of data necessary to perform a thorough reliability analysis. As mentioned earlier, many field data collection methodologies focus on aspects of the field performance other than reliability. Usually, it is a small matter to modify data collection processes to gather the necessary reliability information.

For example, in one instance the field repair personnel were only collecting information specific to the failure of the system and what they did to correct the fault. No information was being collected on the time accumulated on the systems at the time of failure. Fortunately, it was a simple matter to have the service personnel access the usage information, which was stored on a computer chip in the system. This information was then included with the rest of the data collected by the service technician, which allowed for a much greater resolution in the failure times used in the calculation of field reliability. Previously, the failure time was calculated by subtracting the failure date from the date the product was shipped. This could cause problems in that the product could remain unused for months after it was shipped. By adding the relatively small step of requiring the service technicians to record the accumulated use time at failure, a much more accurate model of the field reliability of this unit could be made.

Another difficulty in using field data to perform reliability analyses is that the data may reside in different places, and in very different forms. The field service data, customer support data, and failure analysis data may be in different databases, each of which may be tailored to the specific needs of the group recording the data. The challenge in this case is in developing a method of gathering all of the pertinent data from the various sources and databases and pulling it into one central location where it can easily be processed and analyzed. These functions can also be performed automatically, using ReliaSoft's Dashboard system.

5.2.1 RELIASOFT'S DASHBOARD

ReliaSoft's Dashboard system is a tool for the automation of product quality tracking and warranty processes that pulls in data from a variety of sources and presents the analyzed results in a central location. It is designed around a central database that is used to capture, analyze and present product field reliability, quality and warranty data. The system can be used to capture product quality data for product failures reported via customer returns, the customer call center, field repairs, and other warranty channels. The ReliaSoft Dashboard is a Web-based (or client-server) reporting mechanism that allows users to view product quality and reliability reports and analyses. As needed, customized data entry tools, data load applications, and data transfers from existing systems can be used to capture product data. Following is a description of a Dashboard system currently in use to capture and track field reliability information.

The system is designed around a master database that captures product quality and reliability data into one centralized location. Data are incorporated into the

master database through a variety of techniques designed to work with the information infrastructure already in place at the organization. In this case, Sales, Manufacturing, Warranty Cost, Call Center and Repair Center data are extracted from the databases in which they currently reside on a regular basis and "pumped" into the master database using a specially designed data import utility to validate and load the data. (Other methods to transfer data from existing databases are available and their implementation depends on the existing information architecture and the information technology policies of the individual organization.)

Detailed data on product returns (processed through regional distribution organizations) are captured via a customized Web-based data entry interface that serves to validate and load the data into the master database. Some of the products returned by customers are routed through a product disposition center that has been designed to coordinate more detailed analysis of returned products based on statistical sampling techniques. A customized application in use at this location allows the organization to coordinate the disposition and sampling. More extensive failure analyses are performed on selected products and data obtained during that process is incorporated into the system via a Web-based data entry interface as well. A graphical representation of an example of this process is shown in Figure 2.

Once the pertinent information has been loaded to a central storage area, it can easily be processed and presented in a format that will be meaningful to the users of the information. The Dashboard system presents an at-a-glance overview of a variety of different analyses based on the data that have been extracted. It is then possible to "drill down" to more detailed levels of information from the top-level view.

Another advantage to using the Dashboard (or any other tool or methodology) to pull together disparate sources of field information is that the planning involved in bringing these previously separate sources of information together usually results in a beneficial synergy. By bringing together key players and acting as a catalyst to discussion and discovery, the planning process often helps an organization to gain a new understanding of its current processes and to identify untapped existing resources. Many organizations that have undergone this process have found that the fresh perspective and ability to spark communication provided by such consultations are invaluable to their organizations.

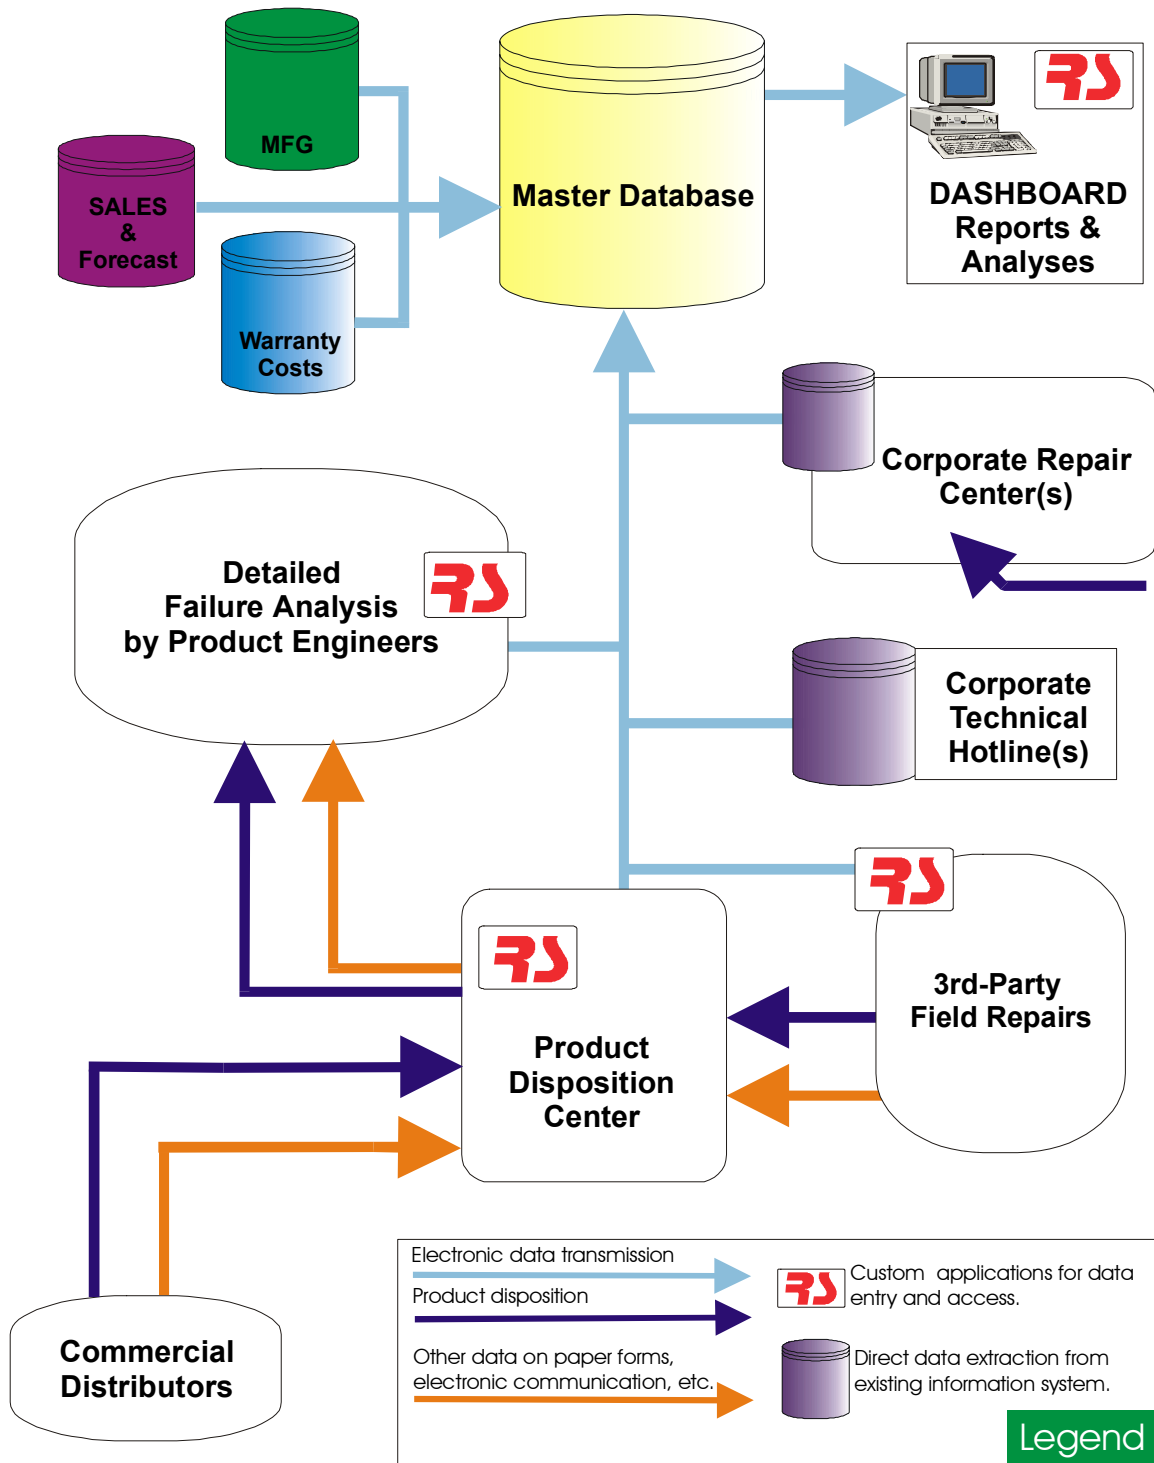


Figure 2 - Example of Linking Existing Data Sources and Customized Data Sources for a Dashboard System

6 DATA ANALYSIS AND REPORTING

The manner in which reliability data is analyzed and reported will largely have to be tailored to the specific circumstance or organization. However, it is possible to break down the general methods of analysis/reporting into two categories: parametric analyses and non-parametric analyses. Overall, it will be necessary to tailor the analysis and reporting methods by the type of data as well as to the intended audience. Managers will generally be more interested in actual data and non-parametric analysis results, while engineers will be more concerned with parametric analysis. Of course this is a rather broad generalization and if the proper training has instilled the organization with an appreciation of the importance of reliability engineering, there should be an interest in all types of reliability reports, at all levels of the organization. Nevertheless, managers are usually more interested in the "big picture" information that non-parametric analyses generally tend to provide, while not being particularly interested in the level of technical detail that parametric analyses provide. On the other hand, engineers and technicians are usually more concerned with the close-up details and technical information that parametric analyses provide. Both of these types of data analysis have a great deal of importance to any given organization, and it is merely necessary to apply the different types in the proper places.

6.1 NON-PARAMETRIC ANALYSIS

Data conducive to non-parametric analysis is information that has not or cannot be rigorously "processed" or analyzed. Usually, it is simply straight reporting of information, or if it has been manipulated, it is usually by simple mathematics, with no complex statistical analysis. In this respect, many types of field data lend themselves to the non-parametric type of analysis and reporting. In general, this type of information will be of most interest to managers as it usually requires no special technical know-how to interpret. Another reason it is of particular interest to managers is that most financial data falls into this category. Despite its relative simplicity, the importance of non-parametric data analysis should not be underestimated. Most of the important decisions that are made concerning the business are based on non-parametric analysis of financial data.

As mentioned in the previous section, ReliaSoft's Dashboard system is a powerful tool for collecting and reporting data. It especially lends itself to non-parametric data analysis and reporting, as it can be quickly processed and manipulated in accordance with the user's wishes.

6.1.1 NON-PARAMETRIC RELIABILITY ANALYSIS

Although many of the non-parametric analyses that can be performed based on field data are very useful for providing a picture of how the products are behaving in the field, not all of this information can be considered "hard-core" reliability data. As was mentioned earlier, many such data types and analyses are just straight reporting of the facts. However, it is possible to develop standard reliability metrics such as product reliability and failure rates from the non-parametric analysis of field data. A common example of this is the "diagonal table" type of analysis that combines shipping and field failure data in order to produce empirical measures of defect rates.

Table 1 gives an example of a "diagonal table" of product shipping and failure data by shipment week. The top row, highlighted in blue and yellow, shows the number of units of product that were shipped in a given week, labeled from 9901 to 9920. The data highlighted in blue and gray represents the number of units

that were reported failed or had warranty hits in the subsequent weeks after being shipped. This information can be used to calculate a simple percent defective for each shipment week. Note that one must make certain to use a weighting factor to account for the amount of time a particular week's worth of units have spent in the field. Also, care should be taken to account for the delay between shipping and actual installation, which can be a substantial time period for some products. The time period for the average delay (in this example two weeks, the data of which appears in the gray diagonal in the table) should be removed from the data being analyzed. Otherwise, a false appearance of a decreasing defect rate appears in the final results of the analysis. Figure 4 shows the results of the non-parametric defect rate calculation, unadjusted and adjusted for the two-week average delay between shipping and installation.

UNITS SHIPPED		530	1290	3362	4071	5966	7516	8174	10235	12578	11932	13236	12965	12559	13052	14153	13705	15122	14389	14400	15306	
		SHIP WEEK																				
		9901	9902	9903	9904	9905	9906	9907	9908	9909	9910	9911	9912	9913	9914	9915	9916	9917	9918	9919	9920	
FAILURE WEEK	9920	2	4	14	14	19	36	28	36	43	39	42	59	46	59	51	55	60	62	26	11	
	9919	2	6	16	13	26	37	36	45	49	50	47	43	49	61	64	56	66	33	16		
	9918	2	5	16	18	28	34	29	38	48	39	65	48	46	49	55	43	26	15			
	9917	2	4	12	17	18	37	38	42	54	56	54	62	52	57	50	24	11				
	9916	2	4	12	16	24	34	38	43	45	40	60	59	48	41	32	16					
	9915	2	5	15	18	24	37	35	39	42	47	58	58	48	22	14						
	9914	2	5	11	20	21	37	35	41	38	43	55	56	23	12							
	9913	2	5	12	20	22	34	33	42	62	52	66	25	14								
	9912	2	5	14	16	24	27	25	42	62	36	22	15									
	9911	2	5	15	14	19	34	37	38	45	22	13										
	9910	2	6	12	18	26	36	38	39	26	11											
	9909	3	5	13	13	20	24	40	23	9												
	9908	2	5	12	17	26	33	16	9													
	9907	2	4	12	13	23	16	7														
	9906	3	6	16	18	12	9															
	9905	2	4	14	9	7																
	9904	2	4	7	3																	
	9903	2	2	3																		
	9902	1	1																			
	9901	0																				
TOTAL FAILURES		39	85	226	257	339	465	435	477	523	435	482	425	326	301	266	194	163	110	42	11	

Table 1 - "Diagonal Table" of Field Data for Non-Parametric Analysis




Figure 3 - Percent Defective Results from Data in Table 1, Unadjusted and Adjusted for Installation Delay

6.2 PARAMETRIC ANALYSIS

Data that lends itself to parametric statistical analysis can produce very detailed information about the behavior of the product based on the process utilized to gather the data. This is the "hard-core" reliability data with all the associated charts, graphs and projections that can be used to predict the behavior of the products in the field.

The origin of this type of data is usually in-house, from reliability testing done in laboratories set up for that specific purpose. For that reason, a great deal more detail will be associated with these data sets than with those that are collected from the field. Unfortunately, in dealing with field data, it is often a matter of taking what you can get, without being able to have much impact on the quality of the data. Of course, setting up a good program for the collection of field data will raise the quality of the field data collected, but generally it will not be nearly as concise or detailed as the data collected in-house.

The exception to this generalization is field data that contains detailed time-of-use information. For example, automotive repairs that have odometer information, aircraft repairs that have associated flight hours or printer repairs that have a related print count can lend themselves to parametric analysis. Caution should be exercised when this type of analysis is performed, however, to make sure that the data are consistent and complete enough to perform a meaningful parametric analysis.

Although it is possible to automate parametric analysis and reporting, care should be taken in automatic processing. Caution is required because of the level of detail inherent in this type of data and the potential "disconnect" between field data and in-house testing data (described in Section 4.1). Presentations of these two types of data should be carefully segregated in order to avoid


unnecessary confusion among the end users of the data reports. It is not unusual for end-users who are not familiar with statistical analyses to become confused and indignant when presented with seemingly contradictory data on a particular product. The tendency in cases such as these is to accuse one or both sources of data (field or in-house) of being inaccurate. This is, of course, not necessarily true. As was discussed earlier, there will usually tend to be a disparity between field data and in-house reliability data.

Another reason for the segregation of the field data and the in-house data is the need for human oversight when performing the calculations. Field data sets tend to undergo relatively simple mathematical processing which can be safely automated without having to worry about whether the analysis type is appropriate for the data being analyzed. However, this can be a concern for in-house data sets that are undergoing a more complicated statistical analysis. This is not to say that parametric analysis should not be in any way automated. However, a degree of human oversight should be included in the process to insure that the data sets are being analyzed in an appropriate manner. Furthermore, the data should be cross-referenced against the Test Log and Service Log to make sure that irrelevant or "outlier" information is not being included in the data.

6.2.1 EXAMPLES OF REPORTING FOR PARAMETRIC DATA ANALYSIS

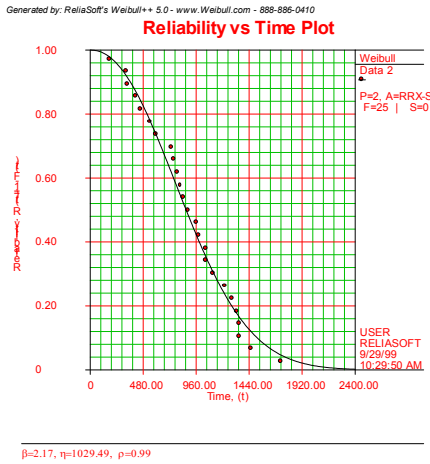
Following are some examples of the information that can be generated using parametric data analysis. While this is by no means complete, it serves as a starting point for the information that can be obtained with the proper collection of data and parametric analysis.

6.2.1.1 PROBABILITY PLOT




Probability plotting was originally a method of graphically estimating distribution parameter values. With the use of computers that can precisely calculate parametric values, the probability plot now serves as a graphical method of assessing the goodness of fit of the data to a chosen distribution. Probability plots have nonlinear scales that will essentially linearize the distribution function, and allow for assessment of whether the data set is a good fit for that particular distribution based on how close the data points come to following the straight line. The y-axis usually shows the unreliability or probability of failure, while the x-axis shows the time or ages of the units. Specific characteristics of the probability plot will change based on the type of distribution..

6.2.1.2 RELIABILITY FUNCTION




The reliability function gives the continuous probability of a successful mission versus the time of the mission. This is similar to the probability plot in that it shows the performance of the product versus the time. However, it does not have nonlinear scales on the axes and the y-axis gives the reliability instead of the unreliability.

6.2.1.3 PROBABILITY DENSITY FUNCTION




The probability density function (*pdf*) represents the relative frequency of failures with respect to time. It basically gives a description of how the entire population from which the data is drawn is spread out over time or usage. The probability density function is most commonly associated with the "bell curve," which is the shape of the *pdf* of the normal or Gaussian distribution.

6.2.1.4 FAILURE RATE FUNCTION




The failure rate function indicates how the number of failures per unit time of the product changes with time. This provides a measure of the instantaneous probability of product failure changes as usage time is accumulated. The failure rate plot is associated with the "bathtub curve," which is an amalgamation of different failure rate curves which illustrates the different ways in which products exhibit failure characteristics over the course of their lifetimes.

6.2.1.5 LIFE VS. STRESS




The Life vs. Stress plot is a product of accelerated life testing or reliability testing that is performed at different stress levels. This indicates how the life performance of the product changes at different stress levels. The gray shaded areas are actually pdf plots for the product at different stress levels. Note that it is difficult to make a complete graphical comparison of the pdf plots due to the logarithmic scale of the y-axis.

6.2.1.6 LIKELIHOOD FUNCTION




The likelihood function is a more esoteric function of the data, but it is directly related to how the parameters are calculated. The likelihood function relates the data points to the values for the parameters of the distribution. The maximization of this function determines the best values for the distribution's parameters.

6.2.1.7 RELIABILITY IMPORTANCE



Reliability importance is a measure of the relative weight of components in a system, with respect to the system's reliability value. The higher the reliability importance a particular component has, the larger the effect that component has on the system's reliability. This measure is useful in helping to optimize system reliability performance, as it helps identify which components will have the greatest effect on the overall system reliability.

6.2.1.8 RELIABILITY GROWTH



Reliability growth is an important component of a reliability engineering program. It essentially models the change in a product's reliability over time and allows for projections on the change in reliability in the future based on past performance. It is useful in tracking performance during development and aids in the allocation of resources. There are a number of different reliability growth models available that are suitable to a variety of data types. The above chart is a graphical representation of the logistic reliability growth model.

7 BRINGING IT ALL TOGETHER

The preceding sections briefly outline some of the basic building blocks of a solid reliability engineering program. These steps are all very helpful in constructing a program that will efficiently gather information, transmit, store, analyze and report on the reliability of an organization's products. The process will not be the same for everyone, of course. The construction or enhancement of a reliability program will by necessity be specially adapted according to the specific needs and structure of the organization. As is the case with many other situations, "form follows function," and the form of the reliability program will follow the function of the organization that is implementing it.

However, it is necessary to make sure that the information that is generated by the reliability program is fed back throughout the organization so that the maximum benefits of the program can be achieved. Instituting a reliability program merely for the sake of having a reliability program will ultimately be of no benefit to anyone. If the reliability program is not feeding back useful information to all of the areas of the organization that need it, it will eventually atrophy and become just a little-utilized enclave of the larger organization. Of course, it is unlikely that an organization that has gone to the trouble of implementing a high-efficiency reliability program will allow such a program to wither and die, but it is important to make sure that the reliability program's benefits reach all the areas that it can.

In the course of this outline, we have discussed some of the more immediate benefits of having a good reliability program in place. Examples include feeding information back to manufacturing organizations to aid in maximizing the efficiency of the manufacturing process and performing system-level reliability analyses that can benefit the early stages of a development program. There are still other methods of putting reliability information to use in order to aid the organization beyond the obvious uses, such as decreased warranty costs.

7.1 CONNECTING FIELD AND LAB DATA

One of the most important activities that can be undertaken once a comprehensive reliability program is in place is to be able to model the transition between reliability data generated as a result of in-house testing and reliability data resulting from the performance of products in the field. The causes for this difference in results as well as the differences in the formats of the data have been discussed elsewhere in this document. Nevertheless, the ability to bridge the difference between these two information sources lies within the grasp of an organization that has a good reliability program and an adequate amount of data. Although it requires a good deal of data manipulation and mathematical analysis, a model can be developed that will allow for the mapping of in-house reliability data to make accurate predictions of field performance. Obviously, this is a powerful tool that would have an important role in projecting warranty costs for new products and the planning of future programs.

7.2 RELIABILITY GROWTH

One use of the information that a reliability program provides is the implementation of a reliability growth study. There are numerous reliability growth models that can be used with a variety of types of input data. The diversity of reliability growth models and acceptable input makes this type of modeling very flexible and it can be applied across a number of different functional areas in an

organization. For example, the detailed data generated during the development phase of a product can be used with a parametric growth model in order to judge whether the project will meet its reliability goal within the allotted time. Based on the growth model results, more efficient allocation of resources on the project could be implemented based on the expected performance of the product. Similarly, a less complicated non-parametric growth model could be used to assess the change of field reliability as a result of design or manufacturing process changes once the product has been released. On a larger scale, the reliability growth of specific product lines can be modeled over the course of several generations of products in order to estimate the reliability and associated warranty costs of future product lines or projects that have yet to be implemented.

7.3 OPTIMUM DESIGN LEVEL DETERMINATION

With a good grasp of the reliability of components and systems, it is possible to devise specifications and designs that result in the optimum level of reliability performance. Designing a product with inexpensive and unreliable parts will result in a product with low initial costs, but high support and warranty costs. On the other hand, over-designing a product with costly, highly reliable parts will result in a final product with low support and warranty costs, but that is prohibitively expensive. Application of information from a reliability engineering program can result in a design that balances out both of these factors, resulting in a design reliability that minimizes the overall cost of the product. Figure 5 gives a graphical representation of this concept.




Figure 5 - Balancing Initial and Post-Production Costs to Determine Optimum Reliability

7.4 COMPETITIVE ASSESSMENT

The principles and practices of reliability engineering that are applied to an organization's products in the normal course of development and production can also be applied to the products of the competition. In setting up a competitive reliability assessment program, a population of a competitor's products is tested in the same manner as those in-house products in development and production. This can provide valuable information as to the relative strengths and weaknesses of the competitors' products. In cases where the reliability or performance of a competitor's product is superior to those being produced by an organization, the competitor's product can be "reverse engineered" in order to gain insight on how the organization's product can be improved. By understanding the performance of the entire gamut of competitive products, an organization can go a long way towards becoming the best in the field.

7.5 MARKETING AND ADVERTISING

In the competitive business world, any edge in helping to find or increase the number of paying customers can result in sizable financial benefits. Given two competing products that are equal in all other respects, the edge belongs to the product that is more reliable. As products become more sophisticated, so do the customers, to the point where the reliability of a product is one of the main considerations a savvy customer takes into account before making a purchase. As a result, more and more advertising includes a reliability slant as part of the sales pitch. From computers to sewing machines, the reliability of the product is increasingly being used to market and sell a variety of products. Some advertisements are now including what were once considered "esoteric" reliability concepts such as MTBF (Mean Time Before Failure) values. With a solid reliability program in place, the information can be used to help sell the product as well as to develop it. This is particularly true if there is data from a competitive assessment program that allows the sales and marketing groups to demonstrate that their product is not only highly reliable but also much more reliable than those of the competition.