

Performing Effective MTBF Comparisons for Data Center Infrastructure

By Wendy Torell
Victor Avelar

White Paper #112

APC[®]
Legendary Reliability[®]

Executive Summary

Mean Time Between Failure (MTBF) is often proposed as a key decision making criterion when comparing data center infrastructure systems. Misleading values are often provided by vendors, leaving the user incapable of making a meaningful comparison. When the variables and assumptions behind the numbers are unknown or are misinterpreted, bad decisions are inevitable. This paper explains how MTBF can be effectively used as one of several factors for specification and selection of systems, by making the assumptions explicit.

Introduction

Avoiding failures in a critical data center is always a top priority. When minutes of downtime can negatively impact the market value of a business, it is crucial that the physical infrastructure supporting this networking environment be reliable. How can one be sure that they are implementing reliable solutions? MTBF is the most common means of comparing reliabilities. However, the business reliability target may not be achieved without a solid understanding of MTBF. The fundamental principles of MTBF are introduced in APC White Paper #78, “Mean Time Between Failure: Explanation and Standards”. Ultimately MTBF is meaningless if the definition of failure is not clear or assumptions are unrealistic or misinterpreted.

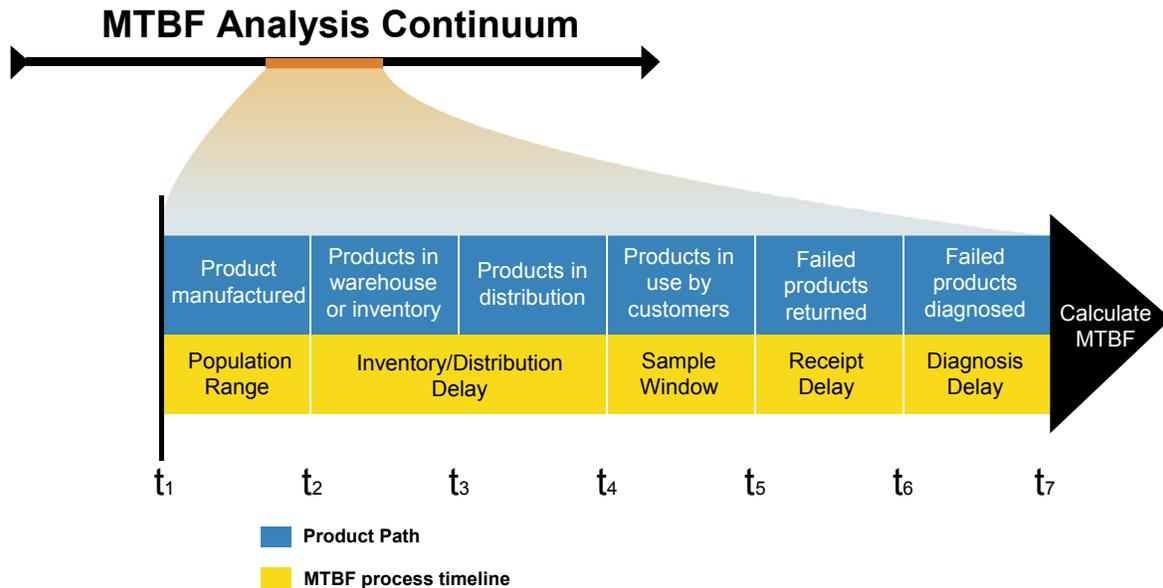
This paper explains how MTBF should be used and the limitations in using it as a specification and selection tool. A checklist is provided as a guideline for ensuring a fair and meaningful cross-system comparison.

Realistic Approach for Comparative MTBF Analysis

In White Paper #78, several methods are introduced for predicting MTBF. With so many methods available, it may appear impossible to find two systems using the same method. There is, however, one method that runs a common thread through the diverse processes of most organizations. The Field Data Measurement Method uses actual field failure data and therefore is a more accurate measure of failure rate than simulations. This data may not be available for products manufactured in low volume or new products, but for products that do have sufficient field populations it should always be used. Therefore it is the most logical and realistic starting point for cross-system comparisons. Note that this method, like many others, is based on the constant failure rate assumption as discussed in White Paper #78.

The steps to the method are introduced in this paper, and the variables within each step that affect the outcome are listed and described. If any of the critical assumptions or variables between systems under comparison vary, it is critical to assess their potential impact on the MTBF estimates. **Figure 1** illustrates the timeline of the field data measurement process. Each of the elements in the timeline is explained in the process steps that follow.

Figure 1 – Field data measurement process



Step 1: Define & estimate size of population

The first step in the process of determining the annual failure rate (AFR), and ultimately the MTBF of a product, is to identify the particular product population to be analyzed. Should the calculation be based off of a particular product model or of an entire product family? How many days or months of manufactured product should be included in the population? When will that production date begin and end? It is important that the product(s) chosen for the population be adequately similar in design and that there be sufficient quantity to have statistical validity for the data gathered.

Step 2: Determine sample time range for collecting data

The second step of the process is to determine the sample time range for collecting the data on failures from the population. The data is often collected when users of the product call the vendor to report a failure. The appropriate amount of time between the population's last manufactured date and the start of the sample period varies with product, geography, distribution process, and inventory position. For example, if units spend two months in the factory warehouse and two months in the distribution pipeline, then the minimum timeframe for the sample period to begin is four months after the close of the population date range. For products that go through distributors, resellers, or retailers, four months is considered a realistic timeframe that accounts for these variables.

There are two important variables here: (1) sufficient time between the population's last manufactured date and the start of the sample period, and (2) a big enough window of data collection to ensure confidence in the results.

If sufficient time is not allowed between the population's last manufactured date and the start of the sample period, then the sample period may begin before the products in the population are fully deployed. Under

this condition two effects happen. First, since units that are not deployed cannot fail, there is a tendency to underestimate the failure rate. The second effect is that the sample period tends to include a large number of installation or setup failures. With new products that may exhibit a failure rate of the classical "bathtub" shape, including a large number of installs causes an overestimation of the failure rate. Although we know that both of these counteracting effects are very strong, we can't say that they balance one another.

The other important consideration with regard to sample time is the duration of the window. How many days are adequate to collect data on failures? The sample time window must be chosen to be wide enough to remove statistical "noise" from the sample. The duration needed to obtain reasonable accuracy is dependent on the size of the population. For instance, this may be one month for a very high volume product, and a few months for lower volume products.

Step 3: Define a failure

Before a failure can be counted it must be well defined to ensure a consistent measurement process. Imagine if a failure was defined by each individual technician as the "failed" products came into the factory. One technician may count only those products that failed catastrophically while another may count all products that failed in any manner including catastrophic. These two extremes would throw off any chance of accurately measuring the failure rate of a particular product. Not to mention the effect it would have on the process control of that product. Therefore, it is imperative that the vendor has a clear definition of failure before diagnosing any products. Sometimes vendors have multiple definitions of failure for calculating the MTBF of specific events. For example, UPS vendors tend to measure the MTBF of products that dropped the critical load as well as less critical failures in which the load continued to operate.

Step 4: Receive, diagnose and repair product

Sufficient time must be allowed between the end of the sample period and the AFR calculation to allow time for products with reported failures to be received, diagnosed, and repaired. The diagnosis determines the type of failure, while the repair validates the diagnosis. For smaller products, the unit is usually sent back to the vendor, which results in a receipt delay or the time required for the unit to arrive. After the unit arrives at the vendor, it must be diagnosed and repaired, which results in another delay called the diagnosis delay. For larger products diagnosis and repairs are usually made at the customer site, therefore, there is little to no delay. In either case, it is necessary for products to be diagnosed and repaired before calculating the AFR. In cases of high volume products it is possible to reach the end of the diagnosis delay and still have units yet to be repaired. In these cases an assumption is sometimes made that unrepaired units fail at the same rate as the previously repaired units. Depending on the manufacturing volume and type of products being measured, receipt delay and diagnosis delay can add weeks to the end of the sample period, at which point the AFR can be calculated.

Step 5: Compute annual failure rate

The annual failure rate is computed to illustrate the expected number of failures in one calendar year of a particular product. The first step in calculating this number is to "annualize" the failure data. This is done by multiplying the number of failures in the sample period by the number of sample periods per year. The

second step is to determine the ratio of failures to the entire population. This is done by dividing the annualized number of failures by the quantity of units built during the population period. **Equation 1** is shown below:

$$AFR = \frac{\text{Failures in the sample period} \times (52 \text{ weeks per year} / \text{Number of weeks in sample period})}{\text{Number of units in population}} \quad \text{Equation 1}$$

This equation makes the following 2 assumptions: (1) the products operate 24 hours a day, 365 days a year; and (2) all products in the population begin operation at the same time. So, even though this formula could be used for any product, it is more relevant for products that are continuously operating. For installations where products are known to run intermittently, it is more accurate to compute AFR using **Equation 2**. One example of this type of product is a standby emergency generator system.

$$AFR = \frac{\text{Failures in the sample period} \times (52 \text{ weeks per year} / \text{Number of weeks in sample period})}{\text{Cumulative operating years of population}} \quad \text{Equation 2}$$

By using this formula, the AFR accounts only for the time that the units are in actual operation. **Equation 1** and **Equation 2** are actually the same equations but with different sets of assumptions. The following hypothetical example illustrates how significant the difference can be when a non-continuously operating product is analyzed:

There are 10,000 automobiles in the sample population.

Over the course of 2 months (sample period), data is collected on failures for this population.

An average automobile operates 400 hours per year.

Throughout the 2 months, 10 automobiles failed.

Using Equation 1:

The failure rate is 10 failures x (52 weeks per year / 8 weeks in sample period) / 10,000 units in population = 0.0065 or 0.65%

Using Equation 2:

Assuming the products went into operation at the same time*, the operational life of the population is 10,000 x 400 hours per year = 4 million cumulative automobile hours or 4 million / 8760 hours per year = 457 automobile years.

The failure rate is 10 failures x (52 weeks per year / 8 weeks in sample period) / 457 cumulative automobile years = 0.14 or 14%

*Note that this assumption was made to simplify the example. In reality, products are sold throughout the period and operating hours decrease as a result. This decrease results in a higher AFR.

If the example above were done with a continuously operating product, the two AFR figures would be identical. Even if the assumption of all units going into operation at the same time were taken out, the AFR numbers would still be fairly close. Therefore, having an understanding as to whether the product will operate continuously or non-continuously is critical to performing a proper analysis.

Step 6: Convert AFR to MTBF

Converting AFR to MTBF (in hours) is the easiest of all the steps but is perhaps the most frequently misinterpreted. Converting AFR to MTBF is valid only under the constant failure rate assumption. The formula is shown below (**Equation 3**):

$$MTBF = \frac{\text{Hours in a year}}{AFR} = \frac{8760}{AFR} \quad \text{Equation 3}$$

Sample MTBF calculation using the AFR measurement process

The following hypothetical example helps to illustrate this entire process.

Step 1:

The population is determined to be all Brand “X” 15kVA UPS systems, manufactured from week 36 to week 47 of 2003 (September 1 through November 21), a 12-week production window. The population consists of 2000 units.

Step 2:

The sample window is determined to begin on February 2, 2004 and end on July 16, 2004 (a window of 24 weeks). This accounts for a 10-week delay for inventory and distribution of the products.

Step 3:

Failures are defined as critical load drops caused by anything, including human error.

Step 4:

During the sample period, twenty failures were reported. Of those, nine were classified as critical load drops and the other eleven were non-critical. So, based on the definition of failure established in Step 3, nine failures are used in the calculation that follows. The failed products were received, diagnosed, and repaired prior to the AFR calculation.

Step 5:

AFR is calculated as:

$$AFR = \frac{9 \text{ failures} \times (52 \text{ weeks per year} / 24 \text{ weeks in sample period})}{2000 \text{ units in population}} = 0.00975 = 0.975\%$$

Step 6:

MTBF is calculated as:

$$MTBF = \frac{8760}{AFR} = \frac{8760}{0.00975} = 898,462 \text{ hours}$$

Variables Affecting AFR

Often times, MTBF values are obtained from vendors without any underlying data to back them up. As mentioned previously, when looking at MTBF figures (or AFR figures) of multiple systems, it is important to understand the underlying assumptions and variables used in the analysis, particularly the way failures are defined. When a comparison is done without this understanding, the risk of a biased comparison becomes high and variations of 500% or more should be expected. This can ultimately lead to unnecessary business

expense and even unexpected downtime. In general, the MTBF values between two or more systems should never be compared without an explicit definition of the variables, assumptions, and definitions of failure. Even if two MTBF values appear similar, there is still the risk of a biased comparison. Therefore it is imperative to look beyond the MTBF results and dissect and understand what goes into those values.

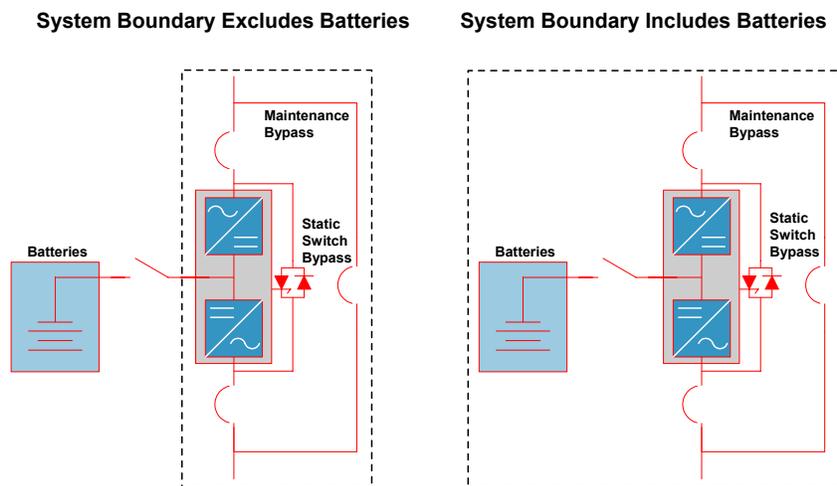
Each variable is explained below and their potential impacts on the results are discussed. As a helpful tool for comparing these variables across two or more systems, a checklist is provided in the appendix. Once complete, the checklist must be reviewed to identify which variables are different across systems. By critically analyzing each of these differences, and their impact on MTBF, it can be determined if a fair comparison is possible as a key input into a product specification or purchase decision.

Product function, application and boundaries

Before comparing two or more MTBF values it is important to verify that the products being compared are equivalent. Products being compared must be similar in function, capabilities, and application. If the product being compared were a UPS, the product function would be to provide back up power to the attached load(s). The application of this product may be to support critical IT loads within a data center environment. Without similar applications, a fair MTBF comparison is not possible. For instance it would be unrealistic to compare one UPS designed for industrial use against one designed for IT use.

More importantly, the boundaries of the systems used in the MTBF comparison must be equivalent. If what is and is not included in each system is set differently, a biased comparison is inevitable. Consider a UPS system with external batteries. Some vendors may choose to exclude any failures resulting from these batteries, since they are “external” and not part of the system. Other vendors may choose to include these battery failures since batteries are an essential component of the system’s operation. **Figure 2** illustrates this example. Other components that might result in inconsistent boundaries include input and output circuit breakers, paralleled systems, fuses, and control systems. Customers should question vendors about what components or subsystems are included in the MTBF calculations and not assume all vendors define things the same way.

Figure 2 – Comparing the “boundaries” for a UPS system



Constant failure rate assumption

For the field data measurement method of computing AFR and MTBF to be valid, the products being analyzed must assume a constant failure rate. It is important to consider whether this assumption is reasonable, given the type of product being compared. This is generally an accepted assumption for electronic systems or components. Do the products fall in this category? If not, the values computed are not likely to be representative of expected failures, which leaves little chance for a fair comparison.

Population size

Once it is clear that the products and their applications are similar, it is important to look at the field data collection process. Defining the population size (number of units produced) is the first critical variable here. If the volume of products defined in the population is too small, the resultant MTBF estimate is likely to be useless. Therefore, when MTBF values are being compared, it is important to make sure each one is based on a sufficient population size.

Although production rates of the products being compared may differ, the important consideration is the number of units in the population. If a product is produced at a lower rate, the timeframe for manufacturing the product should be larger, in order to reach an appropriate volume. For example, vendor "A" produces 1000 units in a month while vendor "B" produces 50 units in a month of an "equivalent" product. Vendor "B" should include several months of manufactured product in their population in order for their result to be statistically valid, while one month should be sufficient for vendor "A".

Time between last manufactured date of population and start of sample period

If sufficient time is not allowed between the end of the population range and the start of the sample collection period, the AFR and MTBF values may be falsely stated. The vendor for each system being compared must provide adequate time for their population to pass through inventory and distribution before beginning the collection of failure data.

For instance, if a particular product is generally in inventory for one month and then goes through a distribution that takes one month, the minimum time that should be allotted before measuring failures is two months. This total "wait" time will vary by product type. Since product types should be similar for a comparison, the time between population and sample periods should be similar. If it is evident that one vendor had insufficient wait time or no wait time at all, their system AFR is likely to appear lower than reality, and caution should be taken in comparing the values.

Sample data collection period

As mentioned in step 2 of the process, it is important that the appropriate sample data collection period be selected. If the systems being compared have the same length sample window with similar production and / or sales volumes, than a fair comparison can be made. However, this may not always be the case. When the length of the collection period varies from one system to the next, it is important to evaluate each one

independently, to determine if they provide an accurate snapshot of the rate of failures that would result over time.

The lower the volume of product, the longer the window should be. For instance, it would not be sufficient for a vendor with a product volume of 10 units per month to collect data on failures for only one month. Because the volume is small, there would be a low degree of certainty that failures (if any) reported in that one-month would project the failure rate over the months ahead.

Definition of a failure

If the failure definition between two comparable products were different, then the analysis would be about as useful as comparing apples to oranges. Therefore, an essential task in carrying out a valid MTBF comparison is to investigate exactly what constitutes a failure for each product being compared. So, what should a vendor consider a failure for the MTBF calculation?

- Is it useful to count failures due to customer misapplication? There may have been human factors that designers overlooked, leading to the propensity for users to misapply the product.
- In the power protection industry, the most popular “definition” of a UPS failure is a “load drop” failure. This means that the power supplied to the load fell outside the acceptable limits and caused the load to turn off. However, is it useful to count load drops caused by a vendor’s service technician? Is it possible that the product design itself increases the failure probability of an already risky procedure?
- If an LED (Light Emitting Diode) on a computer were to fail is it considered a failure even though it hasn’t impacted the operation of the computer?
- Is the expected wear out of a consumable item such as a battery considered a failure if it failed prematurely?
- Are shipping damages considered failures? This could indicate a poor packaging design.
- Are recurring failures counted? In other words, are failures that occur for the same system with the same customer with the same diagnosis counted multiple times or only once?
- Are failures caused during installation counted as failures? This could be the vendor’s technician that caused the failure.
- Are failures counted if the customer did not purchase the recommended maintenance contract or monitoring system?
- If an earthquake results in damage to a building and the system fails, is that counted, or is it excluded as an “Act of God”?
- Are failures of certain components of the system excluded? For a UPS system, this might be the batteries, or the bypass switch.
- If a cascading failure occurs, which brings subsequent systems down, is each system counted as a failure, or just the first?
- If a system is “custom” in some way, does a failure of that system get excluded from the population?

The de facto definition of failure used in the industry to compute MTBF takes into account several deductions. The list above represents just a handful. By making so many exceptions to what counts as a failure, MTBF values present the system as more reliable than what a customer will actually experience. For the purpose of providing partners and customers with AFR and MTBF values, an unambiguous definition of failure is necessary when comparing MTBF values.

Three straightforward definitions are:

- Type 0 The product has a defect or failure that prevents it from being put into operation.
- Type I The termination of the ability of the product as a whole to perform its required function.¹
- Type II The termination of the ability of any individual component to perform its required function but not the termination of the ability of the product as a whole to perform.²

In addition to knowing which definition(s) each vendor has chosen, it is imperative to know whether human causes of failure are included. In cases where human error is included in the MTBF calculation, it becomes much more challenging to compare MTBF numbers. This is because there are many ways in which human error can result in failure, which leads vendors to filter out some of these human-error related failures. If all vendors don't filter out the same types of failures then the system comparison becomes questionable.

To illustrate this point, the Brand "X" example from above will be revisited. **Table 1** compares its MTBF values when different definitions of failure exist. System "A" is the Brand "X" product, where failures are defined as critical (Type I) failures, including all types of human error and failures of consumable items. System "B" is the same Brand "X" product, where failures are also only type I failures, but here they exclude those caused by human error, they exclude cascading failures, and they exclude failures of consumable items. By the nature of the MTBF formula, a difference of even one failure during the sample period can have a significant impact on the MTBF result. In this example, there is a difference of 5 system failures (9 for system A and 4 for system B), and the MTBF varies by 125%. Definitions of failures are easily and often misinterpreted, and as shown in this example, can spell the difference between a valid and invalid comparison. For more information about the tool used to compute the values in this comparison, contact asc@apcc.com.

¹ IEC-50

² IEC-50

Table 1 – Example comparison of MTBF values with different definitions of failure

System A				MTBF Comparison	System B			
Total # of Failures in Sample Period	# of Recurring Failures in Sample Period	# of "1st time" Failures in Sample Period	Include in MTBF Calculation?	Types of Failures	Include in MTBF Calculation?	# of "1st time" Failures in Sample Period	# of Recurring Failures in Sample Period	Total # of Failures in Sample Period
				Type 0 Failure = The product has a defect or failure that prevents it from being put into intended operation				
0		0	<input type="checkbox"/>	Failures from shipping damage	<input type="checkbox"/>	0		0
0		0	<input type="checkbox"/>	Failures caused during a "certified" installation	<input type="checkbox"/>	0		0
0		0	<input type="checkbox"/>	Failures caused during an "uncertified" installation	<input type="checkbox"/>	0		0
				Type I Failure = The termination of the ability of the product as a whole to perform its required function.				
0		0	<input checked="" type="checkbox"/>	"Reported failures" determined to be normal operation	<input checked="" type="checkbox"/>	0		0
1		1	<input checked="" type="checkbox"/>	Cascading failures (i.e. another "like" system has caused this system to fail)	<input type="checkbox"/>	1		1
1		1	<input checked="" type="checkbox"/>	Failures caused by an APC or APC certified service technician (after system is in operation)	<input type="checkbox"/>	1		1
0		0	<input checked="" type="checkbox"/>	Failures caused by a 3rd party technician (after system is in operation)	<input type="checkbox"/>	0		0
1	0	1	<input checked="" type="checkbox"/>	Failures caused by customer misapplication or mis-use	<input type="checkbox"/>	1	0	1
2	0	2	<input checked="" type="checkbox"/>	Failures of consumable items such as batteries	<input type="checkbox"/>	2	0	2
1	0	1	<input checked="" type="checkbox"/>	Hardware component or firmware failures that have since been upgraded or fixed (Engineering Change Orders)	<input checked="" type="checkbox"/>	1	0	1
3	0	3	<input checked="" type="checkbox"/>	* Hardware component or firmware failures	<input checked="" type="checkbox"/>	3	0	3
				Type II failure = The termination of the ability of any individual component to perform its required function but not the termination of the ability of the product as a whole to perform.				
2		2	<input type="checkbox"/>	"Reported failures" determined to be normal operation	<input type="checkbox"/>	2		2
1		1	<input type="checkbox"/>	Cascading failures (i.e. another "like" system has caused this system to fail)	<input type="checkbox"/>	1		1
1		1	<input type="checkbox"/>	Failures caused by an APC or APC certified service technician (after system is in operation)	<input type="checkbox"/>	1		1
1		1	<input type="checkbox"/>	Failures caused by a 3rd party technician (after system is in operation)	<input type="checkbox"/>	1		1
1	0	1	<input type="checkbox"/>	Failures caused by customer misapplication or mis-use	<input type="checkbox"/>	1	0	1
2	0	2	<input type="checkbox"/>	Failures of consumable items such as batteries	<input type="checkbox"/>	2	0	2
1	0	1	<input type="checkbox"/>	Hardware component or firmware failures that have since been upgraded or fixed (Engineering Change Orders)	<input type="checkbox"/>	1	0	1
2	0	2	<input type="checkbox"/>	* Hardware component or firmware failures	<input type="checkbox"/>	2	0	2
Include Recurring Failures in Total?			<input type="checkbox"/>	Recurring Failures are failures for the same customer, same system, and same failure mode	<input type="checkbox"/>	Include Recurring Failures in Total?		
9				Total failures in sample period for MTBF calculation		4		

MTBF Calculation	System A	System B	System B with System A's Definition of Failure
Total Failures in Sample Period for MTBF Calculation	9	4	9
Number of weeks in sample period	24	24	24
Number of units in population	2000	2000	2000
AFR = [Failures in Sample Period x (52 weeks per year / Number of Weeks in Sample Period)] / Number of Units in Population	0.975%	0.433%	0.975%
MTBF = 8760 / AFR	898,462	2,021,538	898,462

System B has a Claimed MTBF 125 % greater than System A. This is an invalid comparison due to differing definitions of failure

System B has an Actual MTBF 0 % greater than System A

In order to alleviate inconsistencies such as this, APC suggests a best practice for defining what is and is not included in an MTBF value. This best practice was established based on the goal of presenting all reasonable failures to customers. These failures should represent all things the vendor has control over. For example, if the vendor's service technician is the cause of a failure, the MTBF should reflect this, since it was the vendor's responsibility. On the other hand, if a customer chose to hire an unauthorized 3rd party service technician, and they were the cause of a failure, the MTBF should not reflect this, since it was out of the vendor's control. The checklist located in the appendix notes which definitions are part of this best practice.

Whenever possible, this best practice definition of failure should be used to compare products across vendors. If a vendor is only able to provide a subset of this definition, then it would be necessary to obtain this same subset from the other vendor being compared. Again, this consistency is necessary to make a fair comparison. However, while this may result in a "fair" comparison, it does not give you a good representation of reality. The smaller the subset of failures included by a vendor, the further from reality the MTBF value becomes.

Time between end of sample period and AFR calculation date

If a vendor could receive, diagnose and repair all product failures reported within the sample period, they could immediately calculate the AFR. In fact, this is possible with lower volume products that are diagnosed and repaired at the customer site. However, this is not the case with higher volume products that are shipped back to the manufacturer. For an MTBF comparison of similar product types, the delay between the end of the sample period and the AFR calculation date should be similar. For example, assume vendor "A" calculates the AFR one month after the close of the sample period and vendor "B" calculates the AFR four months after the sample period. If the product being compared is a high volume product, vendor "A" will most likely report a more favorable AFR. This is because some of their "failed" products (yet to be received, diagnosed and repaired) are not counted in the AFR calculation.

There is one condition where this time range difference between systems is unlikely to result in an invalid comparison (all else being equal). The condition is when all vendors assume that unrepaired units fail at the same rate as previously repaired units and the majority of returns have been received, diagnosed and repaired.

Documented process for data collection and analysis

In order to assess the confidence in an MTBF comparison, it is important to understand the process that each vendor has in place for collecting and analyzing the data. A clearly defined and documented process is critical to implementing a solid quality control program. This helps ensure consistency and accuracy throughout the steps of their analysis. Below are three examples of process problems to look out for. When these or other problems are evident, their impact on the MTBF estimate (and ultimately the comparison) should be closely examined.

- A vendor does not have the ability to track worldwide data with accuracy because different regions of the globe use different tracking systems or storage systems for failure and repair data. Missing or incorrect data can cause errors in the estimation of AFR for units sold internationally.
- A vendor does not have clearly defined processes for categorizing returns. If unused and unopened products that are returned for credit are categorized as returned due to failure, the resulting AFR will be inflated.
- A vendor's tracking system is largely manual. Having more human processes can introduce a range of potential errors in the data and ultimately the AFR calculation. The more automated the process is, the more accurate the results generally are. One example of automation is the scanning of serial numbers instead of the manual typing of numbers into a system.

AFR formula used in calculation

Depending on the product, the AFR formula (**equation 1 or 2**) used by each vendor can render an MTBF comparison useless. Comparing products that operate continuously (once placed in service) can use either formula but comparing products that operate intermittently can use only equation 2, otherwise the comparison is invalid. **Table 2** illustrates under which scenarios a valid comparison can take place.

Table 2 – AFR equation comparison chart

Product operational behavior	AFR equation 1 used	AFR equation 2 used
Continuous operation product comparison I.e. UPS "A" vs. "B" (both backing up critical loads)	Valid comparison	Valid comparison
Intermittent operation product comparison I.e. Laptop "A" vs. Laptop "B"	Invalid comparison	Valid comparison

Hours in a year

Only under the assumption of constant failure rate is it valid to convert AFR to MTBF. In this case, **equation 3** can be used but it is important to verify that all systems in the comparison are using the same number of hours in a year. For example, some vendors use 8,000 hours per year while some use the correct 8,760 hours.

Decision Criterion Beyond MTBF

While MTBF can be a useful decision tool for product specification and selection (when methods, variables, and assumptions are the same for all systems compared), it should never be the sole criterion. There are

many other criteria that should be considered when evaluating products from multiple vendors. For instance, how robust are the vendors' overall quality control processes? What type of volumes are they producing and in what environment? Are they ISO9000 certified? These provide an indication of the standardization of processes for optimizing quality and reliability. How well does each product meet the needs of the user? This may include considerations such as the flexibility or modularity of the product, ability to quickly recover from a failure (MTTR), and the total cost of ownership (TCO) of the product (refer to APC white paper #6, "Determining Total Cost of Ownership for Data Center and Network Room Infrastructure" for discussion on the importance of TCO). Other means of comparing may be to look at customer references or evaluations of the products. Ultimately, an unbiased 3rd party evaluation of the two or more systems under consideration ensures the optimal product specification and purchase decision is made.

Conclusions

When comparing multiple products, MTBF is often a key decision criterion. However, much care should be taken when putting these values side by side. First, the method of predicting the MTBF values must be the same. In addition, many variables and assumptions are used during the process of collecting and analyzing field data and each can have a significant impact on the result. A fair comparison of MTBF is not possible when these variables and assumptions do not line up. The reality is that, often, these variables and assumptions are not the same. The checklist in the appendix can help determine if this is the case. In addition, the online MTBF calculator can help to quantify the impact of critical variables on MTBF values.

With the foundation provided in this paper, MTBF can now be more fairly compared. When similar assumptions and variables are used, and the definitions of failure are the same, there can be a reasonable degree of confidence in the comparison.

About the Authors:

Wendy Torell is an Availability Engineer with APC in W. Kingston, RI. She consults with clients on availability science approaches and design practices to optimize the availability of their data center environments. She received her Bachelors degree in Mechanical Engineering from Union College in Schenectady, NY. Wendy is an ASQ Certified Reliability Engineer.

Victor Avelar is an Availability Engineer for APC. He is responsible for providing availability consulting and analysis for clients' electrical architectures and data center design. Victor received a Bachelor's degree in Mechanical Engineering from Rensselaer Polytechnic Institute in 1995 and is a member of ASHRAE and the American Society for Quality.

Appendix - MTBF Definition of Failure Checklist

Definition of a failure	APC Best Practice	Vendor A	Vendor B
Check off each definition of failure the vendors include in their MTBF values			
Type 0: The product has a defect or failure that prevents it from being put into intended operation			
Failures from shipping damage			
Failures caused during a "certified" installation			
Failures caused during an "uncertified" installation			
Type I: The termination of the ability of the product as a whole to perform its required function			
"Reported failures" determined to be normal operation <i>Two examples of this failure definition are: (1) A UPS switches to battery and drains itself during a blackout thereby dropping the load; (2) An atypical weather condition causes critical servers to shut down because the air conditioning unit could not cool the environment.</i>			
Cascading failures (i.e. another "like" system has caused this system to fail) <i>An example of this failure definition is: There are two paralleled UPS systems on a common output bus. A capacitor on one UPS system shorts causing the fault to propogate to the output bus and drop the load.</i>	√		
Failures caused by an APC or APC certified service technician (after system is in operation)	√		
Failures caused by a 3rd party technician (after system is in operation)			
Failures caused by customer misapplication or mis-use <i>Two examples of this failure definition are: (1) The customer presses the "Off" button instead of the "Test" button causing the load to drop; (2) The customer breaks the chilled water pipes with a forklift thereby causing the air conditioner to stop cooling.</i>			
Failures of consumable items such as batteries <i>Consumable items are defined as any depletable item that should be replaced before the end of a system's useful life. A failure of a consumable item is defined as the termination of the ability of the consumable to perform its expected function prior to the end of its usefull life. Other examples include: (1) Electrolytic capacitors in large systems; (2) Filters such as air and oil filters; (3)The refrigerant inside an air conditioner</i>	√		
Hardware component or firmware failures that have since been upgraded or fixed (Engineering Change Orders) <i>This failure definition includes any Type I hardware or firmware failure that has not been previously counted, which has since been corrected with an ECO or other documented fix.</i>	√		
Hardware component or firmware failures <i>This failure definition includes any Type I hardware or firmware failure that has not been previously counted.</i>	√		
Type II: The termination of the ability of any individual component to perform its required function but not the termination of the ability of the product as a whole to perform			
"Reported failures" determined to be normal operation			
Cascading failures (i.e. another "like" system has caused this system to fail)	√		
Failures caused by an APC or APC certified service technician (after system is in operation)	√		
Failures caused by a 3rd party technician (after system is in operation)			
Failures caused by customer misapplication or mis-use			
Failures of consumable items such as batteries	√		
Hardware component or firmware failures that have since been upgraded or fixed (Engineering Change Orders)	√		
Hardware component or firmware failures	√		