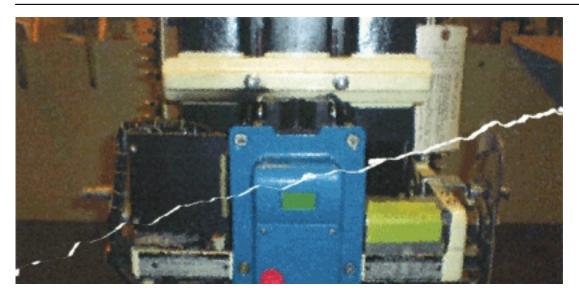
Written by Randall Noon, P.E., Cooper Nuclear Station Wednesday, 01 September 2010 14:06



As this discussion of a power plant's repeated generator-field-breaker failures seems to prove, timing is everything.

In some quarters, if an equipment item fails too often, a frequently chosen *first* solution is to shorten its preventive maintenance (PM) period. This entails either removing the component from service and replacing it with a new one from inventory or perhaps refurbishing the component to a like-new condition and performing a number of tests to verify that its condition meets specifications. This stock solution, however, is not always the best solution: It may actually worsen the item's reliability, and the degraded condition may not be noticed for years. Consider the following real-world account.

Situation

A large power plant uses a motor-generator set to drive a high-pressure, high-volume pump that runs whenever the plant is operating. When the plant reduced power for a scheduled down-power—wherein the pump had to be turned off—the breaker for the field current for the generator failed to open. The circuit for the generator field then had to be opened (at a great inconvenience) farther up-line.

The troublesome breaker had been installed 12 months earlier during the previous outage. It had only been opened and then re-closed one time before. Thus, in terms of operation, this breaker had failed the second time it was opened.

Several breakers of this model are used within the plant. None of them are opened or closed more than once or twice every 18 months. Therefore, operational wear and tear is insignificant.

Written by Randall Noon, P.E., Cooper Nuclear Station Wednesday, 01 September 2010 14:06

A check of maintenance records at both the plant in question and in some similarly designed power-gen facilities found that this model breaker has a checkered history of internal linkage drag and alignment problems. When the breaker fails, it more often does so the first or second time it is opened.

Interestingly, if this type of breaker doesn't fail the first or second time it's opened, it often provides good service for many years before being refurbished: The most commonly reported PM period is 16 years or more. That's why, when this power plant first came online, a PM period of 16 years was established for these units. Due to failures, though, the PM period was at first shortened to nine years, then to 4.5 years. Unfortunately, these changes didn't improve the service reliability of the breaker in question. In fact, a close review of performance data indicated that it was now failing more often. Why didn't the shorter PM improve reliability?

Bathtub curves

Before the answer to the problem can be given, several things need to be explained, including, for example, bathtub curves. A bathtub curve—known formally as a Weibull Distribution Curve—is a plot of the statistical failure rate of a part, component or machine versus time. Fig. 1, which also labels the various parts of the curve, depicts a typical bathtub curve usually shown in textbooks. (This widely used nickname is based on the fact that the curve's statistical distribution plot resembles a cross-sectional view of a bathtub.)

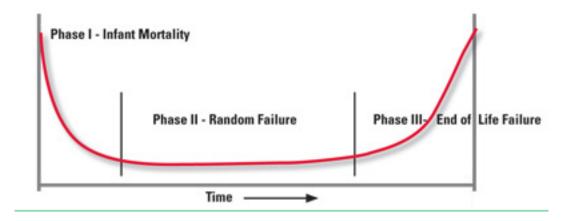


Fig. 1. Failures depicted in a generalized bathtub curve

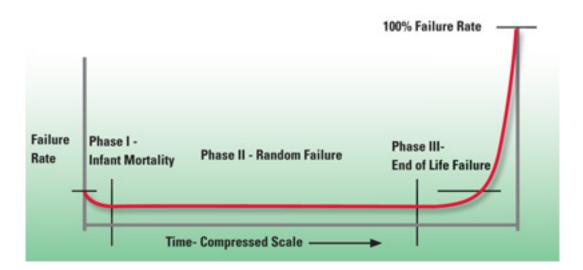
Written by Randall Noon, P.E., Cooper Nuclear Station Wednesday, 01 September 2010 14:06

Phase I, the "Infant Mortality" portion of the curve, reflects the failure rate due to installation errors, assembly errors made at the factory or similar deficiencies. As all the initial problems and "bugs" are found and fixed, the statistical failure rate decreases quickly with time to a minimum level.

Phase II, the "Random Failure" area of the curve, is where the item that has passed through the failures related to assembly error or installation error, operates as designed. Usually this part of the curve is nearly flat and failures are due to statistically random defects and problems.

Phase III, the "End of Life" portion of the curve, is where the item is approaching the end of its useful service life and begins to suffer from wear, age, component breakdowns, environmental degradation and the general curses of entropy.

In depicting the bathtub curve, most textbooks cut the curve off when the Phase III portion is about the same height as the Phase I, "Infant Mortality" portion (as was done below in Fig. 1). Most textbooks also stretch out the time in the Phase I and Phase III portions of the curve, and compress the lapsed time in the Phase II portion. This not only makes the curve appear symmetric (and similar to a bathtub in shape), it makes the curve fit on the page better.



Written by Randall Noon, P.E., Cooper Nuclear Station Wednesday, 01 September 2010 14:06

Fig. 2. Failures depicted in a more realistic bathtub curve

A more realistic bathtub curve is shown in Fig. 2. Note that the Phase III portion of the graph is much higher than the Phase I portion. Statistically, if an item is used long enough, its failure rate, given enough time, approaches 100%, while assembly and installation errors are usually an order of magnitude lower than 100%.

Note in Fig. 2 that the high point of the Phase I region is marked with a horizontal line. An equal failure-rate point is similarly marked with a horizontal line in the Phase III region. When an item has been in service for a long time and is in its end-of-life phase, high reliability is maintained if the replacement or refurbishment period is chosen to be equal to or a little more than the peak infant-mortality failure rate. In other words, don't replace an item if the potential for failure due to installation mess-ups or a factory error is greater than just running the item a while longer. Wait until the chance of failure due to service time equals or exceeds the infant-mortality failure rate.

With this information in mind, let's now re-examine the problem with the breaker.

Revisiting the problem

The conditions cited in the breaker's failure description indicate that the equipment was not failing due to end-of-service-life effects. Parts were not wearing out, nor were they degrading due to age or environment. Furthermore, the facts indicated that if the breaker made it past the first few times it was operated, it would statistically operate in good order for 16 years or more. Thus, the breaker was not experiencing an end-of-life failure; it was experiencing an infant-mortality failure. This distinction is important.

If an item experiences failure due to end-of-life effects and its service life is at the far right on the bathtub curve, shortening the PM period may certainly improve service reliability. As shown in Fig. 2, appropriate shortening of the PM period can shift the PM time from a high failure rate area on the far right to a point on the curve where the failure rate is significantly lower. With some historical data in hand, perhaps the PM can even be shifted to a point where the overall reliability over time is optimized.

Importantly, a shift to the left on the curve only improves reliability if the item is failing due to end-of-service-life effects. In this case, failures were occurring due to infant mortality effects.

Written by Randall Noon, P.E., Cooper Nuclear Station Wednesday, 01 September 2010 14:06

Thus, shortening the PM period actually increased the failure rate. Here is why:

As was learned by checking industry failure statistics for this model breaker, the average refurbishment period is 16 years. This is 3.55 times longer than the 4.5-year PM period the plant was using! If the failure rate due to infant mortality is considered to be P(f), then in 16 years the chance of success with respect to failures caused only by infant mortality factors is [1-P(f)]. If, however, the same breaker is overhauled every 4.5 years, resulting in three refurbishments in that 16-year period, then the expected success rate due to infant mortality is [1-P(f)][1-P(f)][1-P(f)][1].

For example, if the infant mortality rate due to errors in refurbishment or installation is 10%, then the success rate for one breaker for 16 years is 90%. But if the same breaker is refurbished and reinstalled in exactly the same way three times in 16 years, the success rate for that breaker (i.e., no failures occurring) is 73%. As the number of "shots on goal" increases, the chances of a goal being made also increase. Consequently, shortening the PM period for an item that has an infant-mortality problem—such as this breaker—actually decreased overall reliability.

Conclusion

In this case, the solution to the failure problem was NOT to shorten the PM period: Doing so seemed to worsen the failure rate. Since the failure is not a result of end-of-service-life effects, and the characteristics match those of an infant-mortality failure, the solution lies in either fixing the underlying installation or refurbishment deficiency—or changing to a different model breaker with a lower infant-mortality failure rate

MT

Randall Noon is a root-cause team leader at Nebraska's Cooper Nuclear Station. A licensed professional engineer in several states, he's been investigating failures for more than 30 years. Noon is the author of several articles and texts, including: The Engineering Analysis of Fires and Explosions; Forensic Engineering Investigations; and, most recently,

Scientific Method: Applications in Failure Investigation and Forensic Science (published by CRC Press).

E-mail:

rknoon@nppd.com

•