the Availability Digest

www.availabilitydigest.com
@availabilitydig

The Fallacy of Classic Availability Theory

January 2017

In this paper, we point out a fallacy in classic availability theory. We have become accustomed to the terms MTBF and MTR. MTBF, the Mean Time Between Failures, is the average time between failures of a system. MTR, the Mean Time to Repair, is the average time it takes to repair a system. MTBF and MTR are flawed measures, as described below. Instead, we introduce the term MTTF - Mean Time To Failure. This is the expected time to the next system failure. Unlike MTBF and MTR, MTTF is a function of time. As time goes on, MTTF becomes shorter. The likelihood of a system failure draws nearer.

Classic Availability Theory

According to classic availability theory, availability is the proportion of time that the system is operational. Let:

MTBF be the mean (average) time between failures of the system.
MTR be the mean (average) time to repair the system.
A be the probability that the system is operational (it is available).
F be the probability that the system is not operational (it has failed).

Then

$$A = \frac{MTBF - MTR}{MTBF} = 1 - \frac{MTR}{MTBF} = 1 - F \qquad F = \frac{MTR}{MTBF} \qquad (1)$$

Consider a two-node redundant system (either active/passive or active/active) as shown in Figure 1. In an active/backup configuration, one system is acting as the production system; and the other system is standing by to take over in the event that the production system fails. In an active/active system, both systems are actively participating in the application. Should one system fail, all transactions are routed to the surviving system.

The availability of a node, *a*, is a = 1 – MTR/MTBF. The probability that a node will be failed, *f*, is f = (1–a) = MTR/MTBF. The probability that both nodes will be failed (i.e., the system is down), *F*, is

$$F = f^2 = (1\text{-}a)^2 \qquad (2)$$

The probability that the system is up (its availability), *A*, is

$$A = 1 - F = 1 - (1-a)^2 \qquad (3)$$

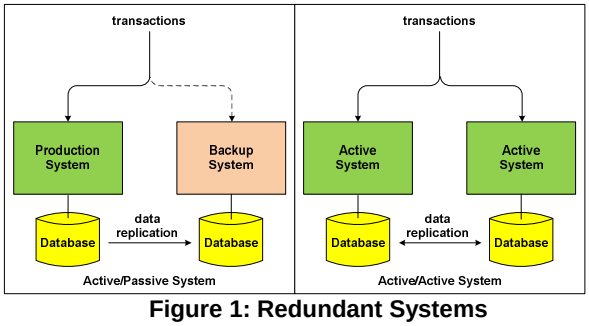

Figure 1: Redundant Systems

© 2017 Sombers Associates, Inc., and W. H. Highleyman
www.availabilitydigest.com
For discussion, contact editor@availabilitydigest.com

Memoryless Variables

In the classic availability theory discussed above, MTBF and MTR are *random variables*. This means that the event (the failure of the system or the repair of the system) is independent of what has happened in the past and has no impact on what will occur in the future. They are *memoryless variables*. This has implications that make no sense:

> Assume that MTBF is 1,000 hours. On the average, failures occur every 1,000 hours. Since MTBF is memoryless, the expected time to the next failure is 1,000 hours. If we wait 500 hours, the average time to the next failure is still 1,000 hours (even if we had a failure in the intervening 500 hours).
>
> Assume MTR is four hours. When the system fails, it will take an average of four hours to repair it. If we wait for two hours and ask the technician when he expects the repair to be completed, he will still say four hours.

The Exponential Distribution

Random variables are described by the exponential distribution function. For instance, consider MTBF. The probability of failure over time is given by

$$p(\text{failure}) = e^{-t/MTBF} / MTBF \qquad 4)$$

The average time to the next failure is

$$\text{average time to next failure} = \int_0^\infty \left(te^{-t/MTBF} / MTBF \right) dt = MTBF \qquad (5)$$

If we wait for a time T, then the average time to the next failure is

$$\text{average time to next failure} = \int_T^\infty \left[(t-T)e^{-(t-T)/MTBF} / MTBF \right] dt = MTBF \qquad (6)$$

The average time to the next failure is still MTBF. Random variables characterized by the exponential distribution function are indeed memoryless.

Classic Availability Theory is Flawed

This is a fundamental flaw in classic availability theory. The time to the next failure is always the same, no matter how long the system has been operating. The time to the completion of the current repair is always the same, no matter how long the system has been under repair.

What is needed is a means to estimate the mean time to the next failure, MTTF, based on realistic probability distributions of failure. MTTF should be a function of time (Figure 2). As time goes on, MTTF should become shorter for realistic systems. It is more likely that the system will fail as time progresses.

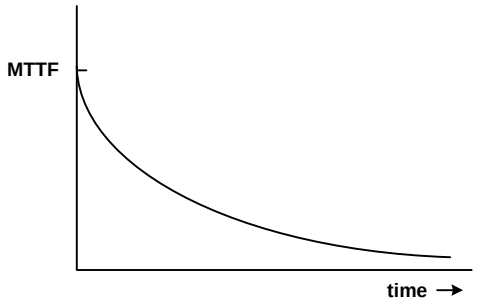

Figure 2: Mean Time to Failure

Mean Time to Failure (MTTF)

Figure 3 shows a typical probability distribution, $p_f(t)$, for the failure of a system. When the system is new, it is unlikely to fail. As it ages, the probability that it will fail increases. At some point, the probability that it will fail will begin to decrease because it likely already has failed.

© 2017 Sombers Associates, Inc., and W. H. Highleyman
www.availabilitydigest.com
For discussion, contact editor@availabilitydigest.com

The probability p_i that the system will fail at some time t_i during a small time interval Δt is $p_i \Delta t$. The mean time to failure, MTTF, for the system is the average of these failure probabilities:

$$MTTF = \sum_{i=0}^{\infty} t_i p_i \Delta t \qquad (7)$$

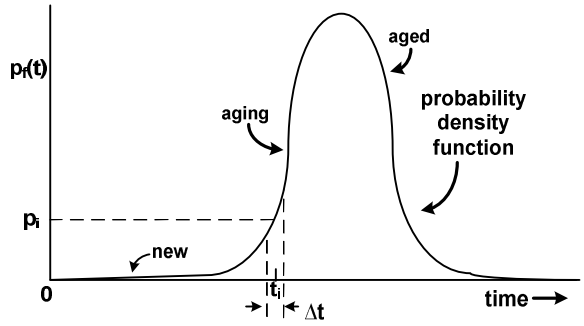

Figure 3: Typical Failure Probability Distribution

For a continuous function, this becomes:

$$MTTF = \int_0^{\infty} t p_f(t) dt \qquad (7)$$

Let us now wait for a time T, as shown in Figure 4. MTTF is now

$$MTTF = \frac{\sum_{i=T}^{\infty} (t_i - T) p_i \Delta t}{\sum_{i=T}^{\infty} p_i \Delta t} = \frac{\sum_{i=T}^{\infty} t_i p_i \Delta t}{\sum_{i=T}^{\infty} p_i \Delta t} - T \qquad (9)$$

where the MTTF term has been normalized to account for the shorter time. Comparing equations (7) and (9), it is clear that MTTF has become shorter as time has progressed (except for the unusual case in which the system survives into old age).

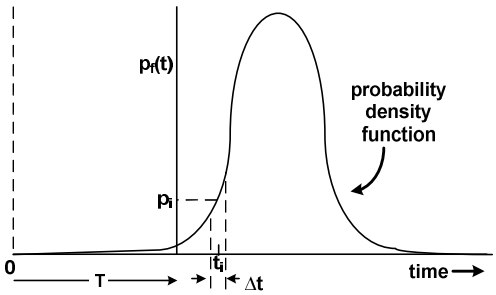

Figure 4: Failure Probability at a Later Time

© 2017 Sombers Associates, Inc., and W. H. Highleyman
www.availabilitydigest.com
For discussion, contact editor@availabilitydigest.com

Redundant System

As described earlier, the reliability of a system can be greatly improved by making it redundant. A second system is added. As shown in Figure 1, the redundant pair can be operated either as an active/backup pair or as an active/active system.

Figure 5 shows a typical system failure probability distribution including infant mortality. Infant mortality is a system failure caused by defects not found in its initial testing before installation. In some cases, the system may not come up at all. In other cases, it may fail shortly after it becomes operational. Once the system is "burned in," the system will run reliably until it ages.

A redundant system is available so long as one of the systems is operational. It fails only if both systems fail.

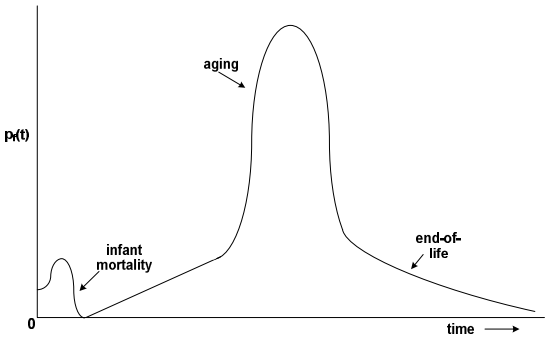

Figure 5: Infant Mortality

In a dually redundant system comprising a System 1 and a System 2, let the probability distribution of failure for System 1 be $p_{f1}(t)$ and the probability distribution of failure for System 2 be $p_{f2}(t)$. The mean time to repair a system is MTR. The probability distribution that both systems will fail is MTR $p_{f1}(t)\ p_{f2}(t)$, as shown in Figure 6. Clearly, the probability that both systems will fail simultaneously is less than the probability that either system will fail at that time. The peak probability that both systems will fail occurs at the peak probability of each failure probability distribution.

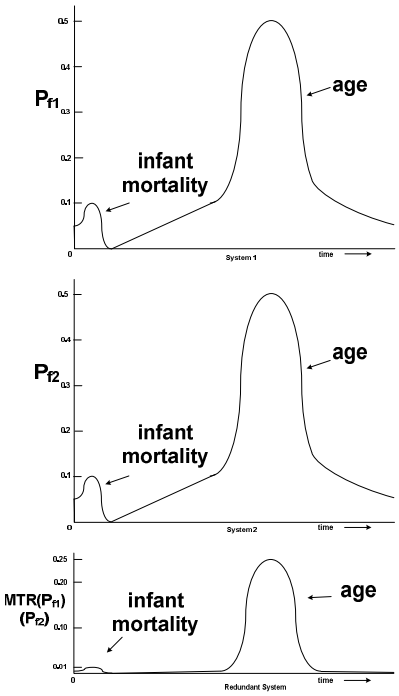

Figure 6: Systems Started Simultaneously

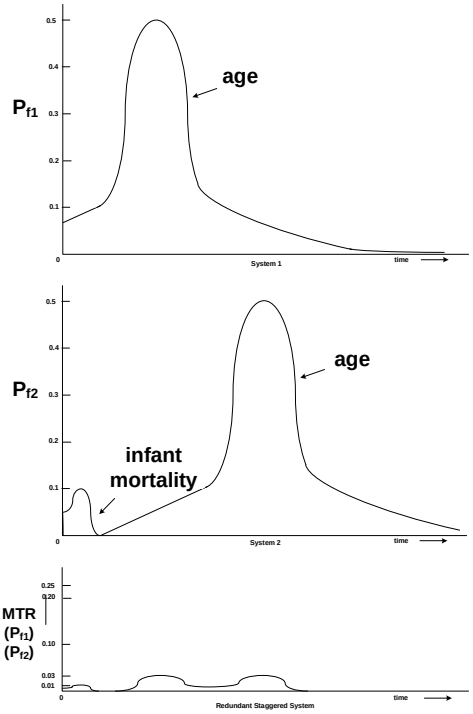

Figure 7: Starting Times Staggered

© 2017 Sombers Associates, Inc., and W. H. Highleyman
www.availabilitydigest.com
For discussion, contact editor@availabilitydigest.com

The availability of the redundant system can be significantly improved by staggering the starting times of the two nodes, as shown in Figure 7. When the probability of failure of one system is high, the probability of the other system is low, thus minimizing the chance that there will be a dual system failure.

Summary

Classic availability theory is flawed in that the expected time to a system failure does not change with time. Clearly, as time goes on, the expected time to system failure should shorten. This flaw is corrected with the concept of Mean Time to Failure (MTTF). MTTF can be used to determine the impact on the availability of various redundant system configurations.

Acknowledgement

Dr. Bruce D. Holenstein, CEO and President of Gravic, Inc., envisioned the effect of staggering systems. Our thanks to him for this insight.

© 2017 Sombers Associates, Inc., and W. H. Highleyman
www.availabilitydigest.com
For discussion, contact editor@availabilitydigest.com