

Increasing Availability at the Cost of Reliability

By Dan Byron

A user considers a system reliable if it is available and operational when needed. One method of increasing system availability is to add more components to the system, but increasing the number of components increases the failure rate of the system—thereby decreasing reliability from an engineering standpoint. This article reconciles the apparent disparity between reliability and availability, and examines redundancy as a way to increase availability.

From an engineering standpoint, reliability is the ability of a system or unit to perform a required function (continue to operate) under stated conditions (laptop unit, IT center) for a specified period of time (three years, as an example). From a customer's point of view, a system is reliable if it operates properly each time the customer wishes to use it. This operational reliability may be more aptly defined as availability: the system is available and fit for use whenever required.

A highly available system need not be highly reliable, although this situation is ideal. An examination of how reliability is assessed provides a basis for understanding the relationship between reliability and availability.

MTBF as a measure of reliability

Each assembly, subassembly, device, or component within a system has its own inherent reliability, often expressed as a mean time between failure (MTBF). The inherent reliability of a system is a function of the sum of the non-reliabilities (failure rates) of all components in the system.

Consider an integrated circuit (IC) that has an MTBF of 100,000 hours.¹ If this device is placed into a circuit that also contains an LED with an MTBF of 100,000 hours, the reliability of the circuit thus far is not additive; it is not 200,000 hours. To determine the MTBF of the circuit, first convert the MTBF of each component into its corresponding failure rate (the reciprocal of its MTBF):

$$\begin{aligned} 100,000 \text{ hour MTBF of IC} &= 1/100,000 = 0.00001 \\ 100,000 \text{ hour MTBF of LED} &= 1/100,000 = \underline{0.00001} \\ \text{Sum} &= 0.00002 \end{aligned}$$

Take the reciprocal of the sum to produce the MTBF of the system (circuit):

$$= 1/0.0002 = 50,000 \text{ hours}$$

This methodology can be applied to any system that is serial in nature: the input of one element depends on the output of another, and the failure of any device will produce a failure of the entire system.

Intuitively, the more material that is added to a serial system, the lower the resultant MTBF, or the higher the failure rate, will be. This premise is the basis for seeking alternate design approaches, such as fault tolerance, fault resilience, or redundancy, which try to keep a system operational even in the event of a hardware failure—satisfying the customer's definition of reliability.

The impact of usage and averages

Consider a commercial desktop unit in which a design team plans to use a power supply with a calculated reliability of 10,000 hours MTBF. Before placing this power supply in a customer's system, the design team needs to consider two additional factors: the customer's usage and the nature of the reliability prediction.

¹ The specifications used in this and other examples in this article are not based on actual data unless otherwise noted.

A typical customer may operate a desktop machine 5 days a week, 50 weeks a year, 12 hours per day: $5 \times 50 \times 12 = 3,000$ hours of intended usage per year. If the machine has a warranty or hardware refresh period of three years, then a power supply with an MTBF of 10,000 hours seems adequate— $3 \text{ years} \times 3,000 \text{ hours/year} = 9,000$ usage hours, which is less than the 10,000 hour MTBF of the power supply.

This assumption is problematic because of the nature of the reliability data itself. MTBF is an *average* value of a universal population of like devices. Although the power supply has a 10,000 hour MTBF, some supplies will fail in the first few hours and some may not fail for 100,000 hours or more. However, in a universal population, these power supplies on average will fail every 10,000 hours.

Another problem with the 10,000 hour MTBF figure is that it represents machine operational hours, not calendar time. That is, 10,000 hours can be accumulated by one machine operated for slightly more than three years, as in the previous scenario, or by 10,000 machines operated for one hour each.

Consider a corporation with 2,000 employees, each using a system that contains the power supply with a 10,000 hour MTBF. This type of power supply will fail at a rate of 0.0001 supplies per hour. Two thousand machines operating per hour multiplied by the failure rate of 0.0001 results in 0.2 failures per hour. To calculate the failure rate during a 12-hour day, multiply the number of hours by the hourly failure rate: $12 \text{ hours per day} \times 0.2 \text{ failures per hour} = 2.4$ failure events per day at this particular corporate site.

The exponential effect of time

Exacerbating these issues is the fact that the reliability of a unit changes with time. The following equation, which forms the basis of nearly every reliability function, indicates that the reliability of an electrical unit exponentially decreases through time:

$$R = e^{-(t/MTBF)}$$

where

- R = the probability that the unit will be fit for use
- e = the natural log
- t = the time under consideration
- MTBF = the reciprocal of the failure rate

As *t* changes, the residual reliability *R* changes. For example, the reliability of a power supply with a 10,000 hour MTBF is 74.08 percent at the end of the first year of operation; that is, 25.92 percent of the population may have already failed:

$$R = e^{-(3000/10000)} = e^{-0.3} = 0.7408$$

At the end of the second year, the residual reliability of the power supply is 54.88 percent and decreases to 40.65 percent at the end of the third year.

Improving availability through redundancy

The high failure rate may distress the customer because each power supply failure renders a desktop unit unavailable. So that the unit can tolerate a failure without a loss of usability (increased availability), the design team considers adding more hardware (reducing reliability by increasing the failure rate of the system).

For this customer, the design team decides to use redundant power supplies. They will configure two power supplies so that one provides all of the power for the system when the other fails.

In this example, assume that the power supplies are load-sharing and that any switching logic, handshaking, error reporting, and so forth are flawless. The only remaining task is to determine how much better, from the perspective of availability, this arrangement of supplies will be over a single-supply configuration.

Reliability of load-sharing devices

The failure rate of a load-sharing element increases because of the added stress that occurs when the first half of the pair fails. Consider the case of load-sharing power supplies: When both are present and operational, each contributes only 50 percent of the system’s power. When the first failure occurs, the remaining supply must increase its output to provide 100 percent of the power.

A component in the power supply that is rated to withstand 1 ampere (A) of current will have a lower failure rate if it is subjected to only 0.5 A, which should occur when both supplies are operational. However, when the first supply fails, the component in the second supply will be subjected to 1 A and will need to operate at this level until the failed supply is replaced. (Component derating guidelines, such as those adopted by Dell® Design Engineering, would prevent a component rated at 1 A from being used in a circuit that may be subjected to 1 A.)

After the first failure, when the customer’s power system is no longer redundant, the reliability of the overall system can be calculated using Markov analysis. The system will operate in a degraded state until the failed supply can be replaced or repaired.

From the Markov model perspective (Figure 1), the overall reliability of the system changes when a supply fails and when it is replaced. Availability is affected when the first supply fails because the system is no longer protected against a power supply failure. Availability returns to higher levels when the failed supply is replaced and when each of the two supplies resumes contributing 50 percent of the system’s power.

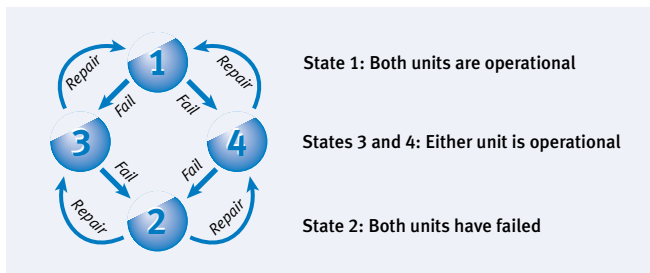


Figure 1. State diagram of two load-sharing power supplies

Year	Non-redundant	Load-sharing
1	74.08	99.95
2	54.88	99.90
3	40.65	99.85

Figure 2. Residual reliability of power supply systems

The reliability of the power supply system increases significantly from 10,000 hours per supply to 6.26 million hours for the load-sharing pair. This number represents the time between system outages caused by the failure of the second supply prior to the repair or replacement of the first. The reliability calculation presumes that an administrator receives failure notification at the point of first failure and completes the repair within eight hours.

Impact of improved availability on failure rate

From an availability perspective, the reliability of the dual-supply power system improves dramatically (see Figure 2) because of the redundancy of two power supplies in the desktop unit. However, the design team must consider the cost in terms of failure rate.

An earlier example showed that 0.2 power supplies per operating hour were likely to fail when each system contained one power supply with an MTBF of 10,000 hours. In a redundant system, twice as many supplies will fail because the system contains twice as many units. However, as long as a failed supply can be identified, isolated, removed, and repaired or replaced before the remaining supply fails, the system will continue to operate (be available) and thus meet the customer’s definition of availability.

Reliability of power supplies in the PowerEdge 2550

To provide a “real world” example, the reliability of a Dell PowerEdge® 2550 will be calculated. This product was selected for two reasons: it has been in the general population long enough to accumulate significant hours of credible data, and it can operate in a single or dual power supply configuration.

The power supply was designed with a specified MTBF of 400,000 hours. Although the true MTBF of the power supply in the field is well above this value, this example will use the specified MTBF.

In the field, the PowerEdge 2550 has a population exceeding 35,000 units; for simplicity, consider this number to be the universal population. Further assume that each PowerEdge 2550 operates in a 24x7 environment.

The hours of accumulated power supply or system usage are 24 hours/day x 365 days/year x 35,000 = 306,600,000 hours/year. The power supply and system values are the same only when the units operate in a single power supply configuration.

Given the MTBF specification of 400,000 hours, or the failure rate of 0.0000025 (reciprocal of the MTBF), the number of expected failures can be calculated as either 306,600,000/400,000 = 766 or 306,600,000 x 0.0000025 = 766. That is, in a single power supply configuration where a power supply failure will render the

PowerEdge 2550 inoperable, 766 units in a universal population of 35,000 would be expected to suffer such a fate in the first year of operation. Each power supply with an MTBF of 400,000 hours that has operated for one year has a residual reliability of 97.83 percent.

For many customers, this potential risk of failure is too high. To mitigate the risk, redundant power supplies are an option.

In the case of redundant supplies, the residual reliability at the end of one year increases from 97.83 percent to 99.9997 percent, and the MTBF (viewed as an availability) increases from 200,000 hours (two supplies at 400,000 hours each) to 10 billion hours. That is, although each supply in the pair is expected to fail every 400,000 hours (on average in a universal population), at least one of the two supplies will be available for 10 billion hours.

The MTBF increases to 10 billion only if the system (software, background diagnostics, wellness monitoring activity, and so forth) notifies an operator that an error is about to occur and if the operator can replace the power supply in a non-disruptive manner. If the failure goes undetected or the system does not produce any failure notifications, then the system does not benefit from redundancy: the failure of the first supply returns the reliability model to that of a serial system, and a second power supply failure will render the system unusable.

Given the increased MTBF of the redundant pair, the number of expected system outages is reduced from 766 occurrences to less than 1 (306,600,000/10,000,000,000 = 0.03). However, this increased availability comes at a cost in terms of spares, logistics, staffing, and related phenomenon because the number of power supplies in the population has doubled. Rather than replacing potentially 766 failed supplies in the first year, an administrator may need to replace 1,533 supplies—without affecting the business—because each of the 35,000 systems contains twice as many power supplies.

Increased reliability in the field

The PowerEdge 2550 and other Dell products operate in excess of calculated reliability levels because of the conservative nature of reliability calculations and the effective, proactive use of Dell tools and software. For example, a PowerEdge 2550 with dual power supplies would have an expected reliability of 99.9997 percent if calculated using the reliability of the power supplies alone. This number would be reduced when the remaining assemblies and subassemblies are added to the system. However, when the full configuration of the PowerEdge 2550 (power supplies, fans, control cards, disks, and so on) is considered, the typical availability of the system in the field exceeds 99.99 percent. ☺

Dan Byron (dan_byron@dell.com) is a senior reliability consultant on the Reliability Engineering Storage Products team at Dell. Prior to joining Dell, he was the reliability engineering manager at EMC Corporation. Dan has a B.A. in Management from the University of New Hampshire and continuing education units in Applied Reliability Engineering from the University of Arizona.