

White paper: Measuring high availability power protection systems: the power availability index

Power availability relies on UPS redundancy, and hot-swappability

High availability is one of the most important issues in computing today. Understanding how to achieve the highest possible availability of systems has been a critical issue in mainframe computing for many years, and now it is just as important for IT and networking managers of distributed processing. A certain amount of mystery surrounds the topic of power availability, but consideration of just a few important points leads to a metric which IT managers can use to increase their systems and applications availability and make a rational price/performance purchase decision.

The importance of high systems availability

Availability is a measure of how much time per year a system is up and available. Usually, companies measure application availability because this is a direct measure of their employees' productivity. With critical applications, or parts of critical applications, physically distributed throughout the enterprise, and even to customer and supplier locations, IT managers need to take the necessary steps to achieve high applications availability throughout the enterprise.

Power availability is the largest single component of systems availability and is a measure of how much time per year a computer system has acceptable power. Without power, the system, and most likely the application, will not work. Since power problems are the largest single cause of computer downtime, increasing power availability is the most effective way for IT managers to increase their overall systems availability. Power availability, like both systems and applications availability, has two components: mean time between failures (MTBF) and mean time to repair (MTTR). The two most important issues in increasing power availability are therefore increasing the MTBF and decreasing the MTTR of the power protection system.

Increasing MTBF

MTBF is the average number of hours it takes for the power protection system to fail. The MTBF of the system can be increased in two ways: by increasing the reliability of every component in the system, or by ensuring that the system remains available even during the failure of an individual component. There is a finite limit to how reliable individual components can get, even with increased cost. Today, typical power protection systems that rely only on high component reliability achieve MTBF between 50 000 hours and 200 000 hours.

By adding a level of redundancy to the system it is possible to achieve a three- to six-fold improvement in MTBF for power protection devices. Redundancy means that a single component of a power protection system can fail and the overall system will remain available and protect the critical load.

Of course, component reliability is a requirement of any system. However, Figure 1 shows the diminishing returns of increasing component reliability. Line 1 shows the plateau that occurs when MTBF is increased by using more reliable (and therefore more costly) components. Line 2 shows how redundancy, in addition to component reliability, can raise MTBF to the next plateau.

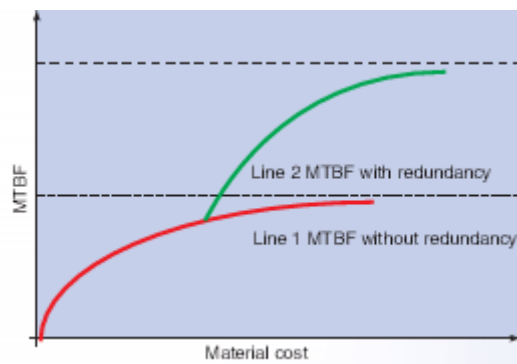


Figure 1 The diminishing returns of component reliability

Decreasing MTTR

One way that systems downtime can occur is when both the power protection system and the utility power fails. A shorter MTTR can decrease the risk that both of these events will occur at the same time. By driving the MTTR towards zero, it is possible to essentially eliminate this failure mode.

Adding hot-swappability to a power protection system is the most effective way of decreasing MTTR. Hot-swappability means that if a single component fails, it can be removed and replaced by the user while the system is up and running. When hot-swappability is used in conjunction with a redundant system, MTTR is driven close to zero, since the device is repaired when there is a component failure but before there is a systems failure.

The Power Availability (PA) Chart

The relationship between power availability, redundancy, and hot-swappability is easily explained by using the PA Chart, which categorises power protection systems in quadrants according to how well they meet the requirements of high power availability – redundancy and hot-swappability. As more components in a system become hot-swappable, the system moves from the bottom to the top of the graph (Figure 2), and as more components become redundant, it moves from the left to the right of the graph. IT managers can choose the solution that is right for them, depending on the need for high availability and the amount of money they want to spend.

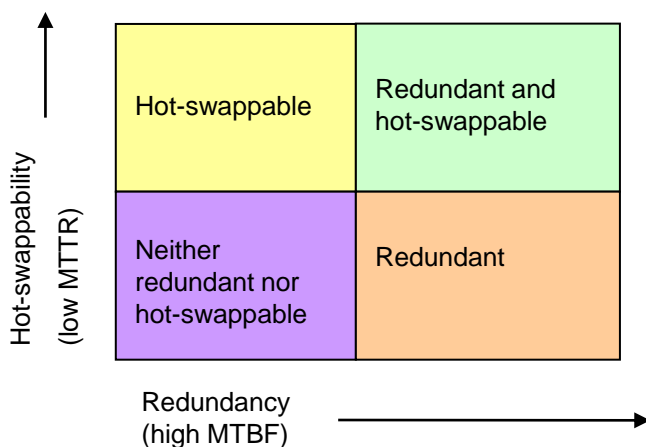


Figure 2 Power protection systems can be categorised according to how well they meet the requirement of high power availability



The PA Chart corresponds to the types of power protection systems available today as shown in Figure 3. The standalone UPS is neither hot-swappable nor redundant. As shown in the table, a standalone UPS provides normal power availability because uptime is dependent on the reliability of the UPS itself.

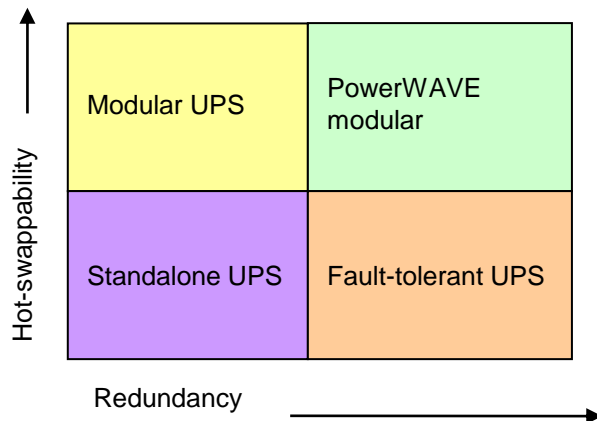


Figure 3 Different types of UPS mapped onto the PA Chart

The fault tolerant UPS is sometimes described as providing affordable redundancy. Systems of this type have redundant components but not all of the major components are hot-swappable. This type of system offers high power availability because the power protection system will continue to protect the load when a component fails. But because a failed component often results in the entire UPS needing replacement, this type of system can have serious drawbacks, including expensive and time-consuming repair with both systems downtime and a major inconvenience for IT managers. Fault tolerant UPS systems may have some hot-swappable components, such as batteries and a subset of power electronics, but in most cases a high number of critical components, such as the processor electronics, will not be hot-swappable. The more components that are not hot-swappable, the lower the power availability.

Like fault-tolerant UPS, modular UPS offer high power availability. Modular UPS have multiple hot-swappable components and are typically used for multiple servers and critical applications equipment. Many modular UPS also have redundant batteries. Their main advantage over fault-tolerant UPS is that all of the main components which can potentially fail can be hot-swapped, eliminating planned downtime due to a service call.

The PowerWAVE range of modular UPS offers the highest level of power protection currently available in the UPS market. In a PowerWAVE modular UPS the power electronics, batteries, and processor electronics are both redundant and hot-swappable. This system provides very high power availability and the highest level of protection for IT managers' critical loads. A PowerWAVE modular UPS may cost a little more than a similarly-rated standalone UPS, but the increased system reliability and availability are invaluable to the IT manager.

The Power Availability (PA) Index

The different types of power protection systems in the PA Chart can be measured linearly with the PA Index, according to how much power availability they provide. The PA Index serves as a tool to explain the difference between power protection systems. Figure 4 shows each of the quadrants from the PA Chart mapped into a level of the PA Index.



Power availability index	Definition	Power availability
PA-1	Not hot-swappable and not redundant	normal
PA-2	Redundant but not hot-swappable	high
PA-3	Hot-swappable but not redundant	high
PA-4	All main components are redundant and hot-swappable	very high

Figure 4 The quadrants of the PA Chart mapped into a level of PA Index shows each of the quadrants from the PA Chart mapped into a level of the PA Index.

Figure 5 shows the relative power availability provided by each type of system. The PA Index maps directly into the PA Chart and makes the different characteristics of high availability power protection systems clear.

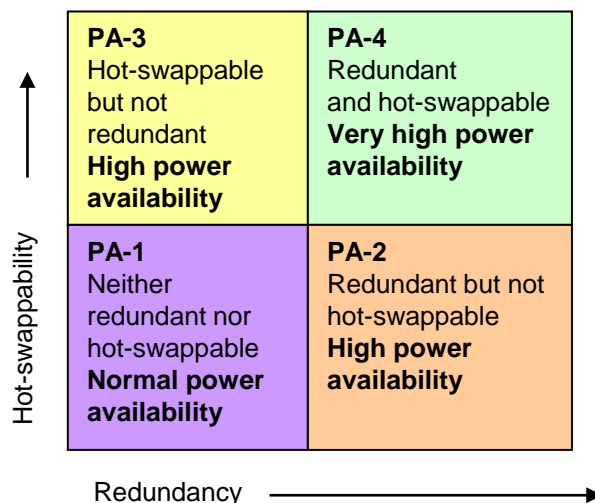


Figure 5 The PA Index mapped into the PA Chart clarifies the relationship between power protection system characteristics and power availability

Conclusion

In conclusion, IT managers can use the PA Chart and the PA Index to help them choose the right power protection system for their high availability applications. The standalone UPS, the modular UPS, and the PowerWAVE 9000DPA Series modular UPS all offer real benefits in terms of power availability versus cost. Although fault-tolerant UPS offer high power availability – and are marketed as such – they introduce serious drawbacks including a high MTTR and potentially significant inconveniences for IT managers.

Contact

Uninterruptible Power Supplies Ltd
 Woodgate
 Bartley Wood Business Park
 Hook
 Hampshire
 RG27 9XA

Phone: 01256 386700
 Email: sales@upspower.co.uk
 Web: www.upspower.co.uk