WHITE PAPER

How to optimise your data centre's uptime and how your UPS can contribute

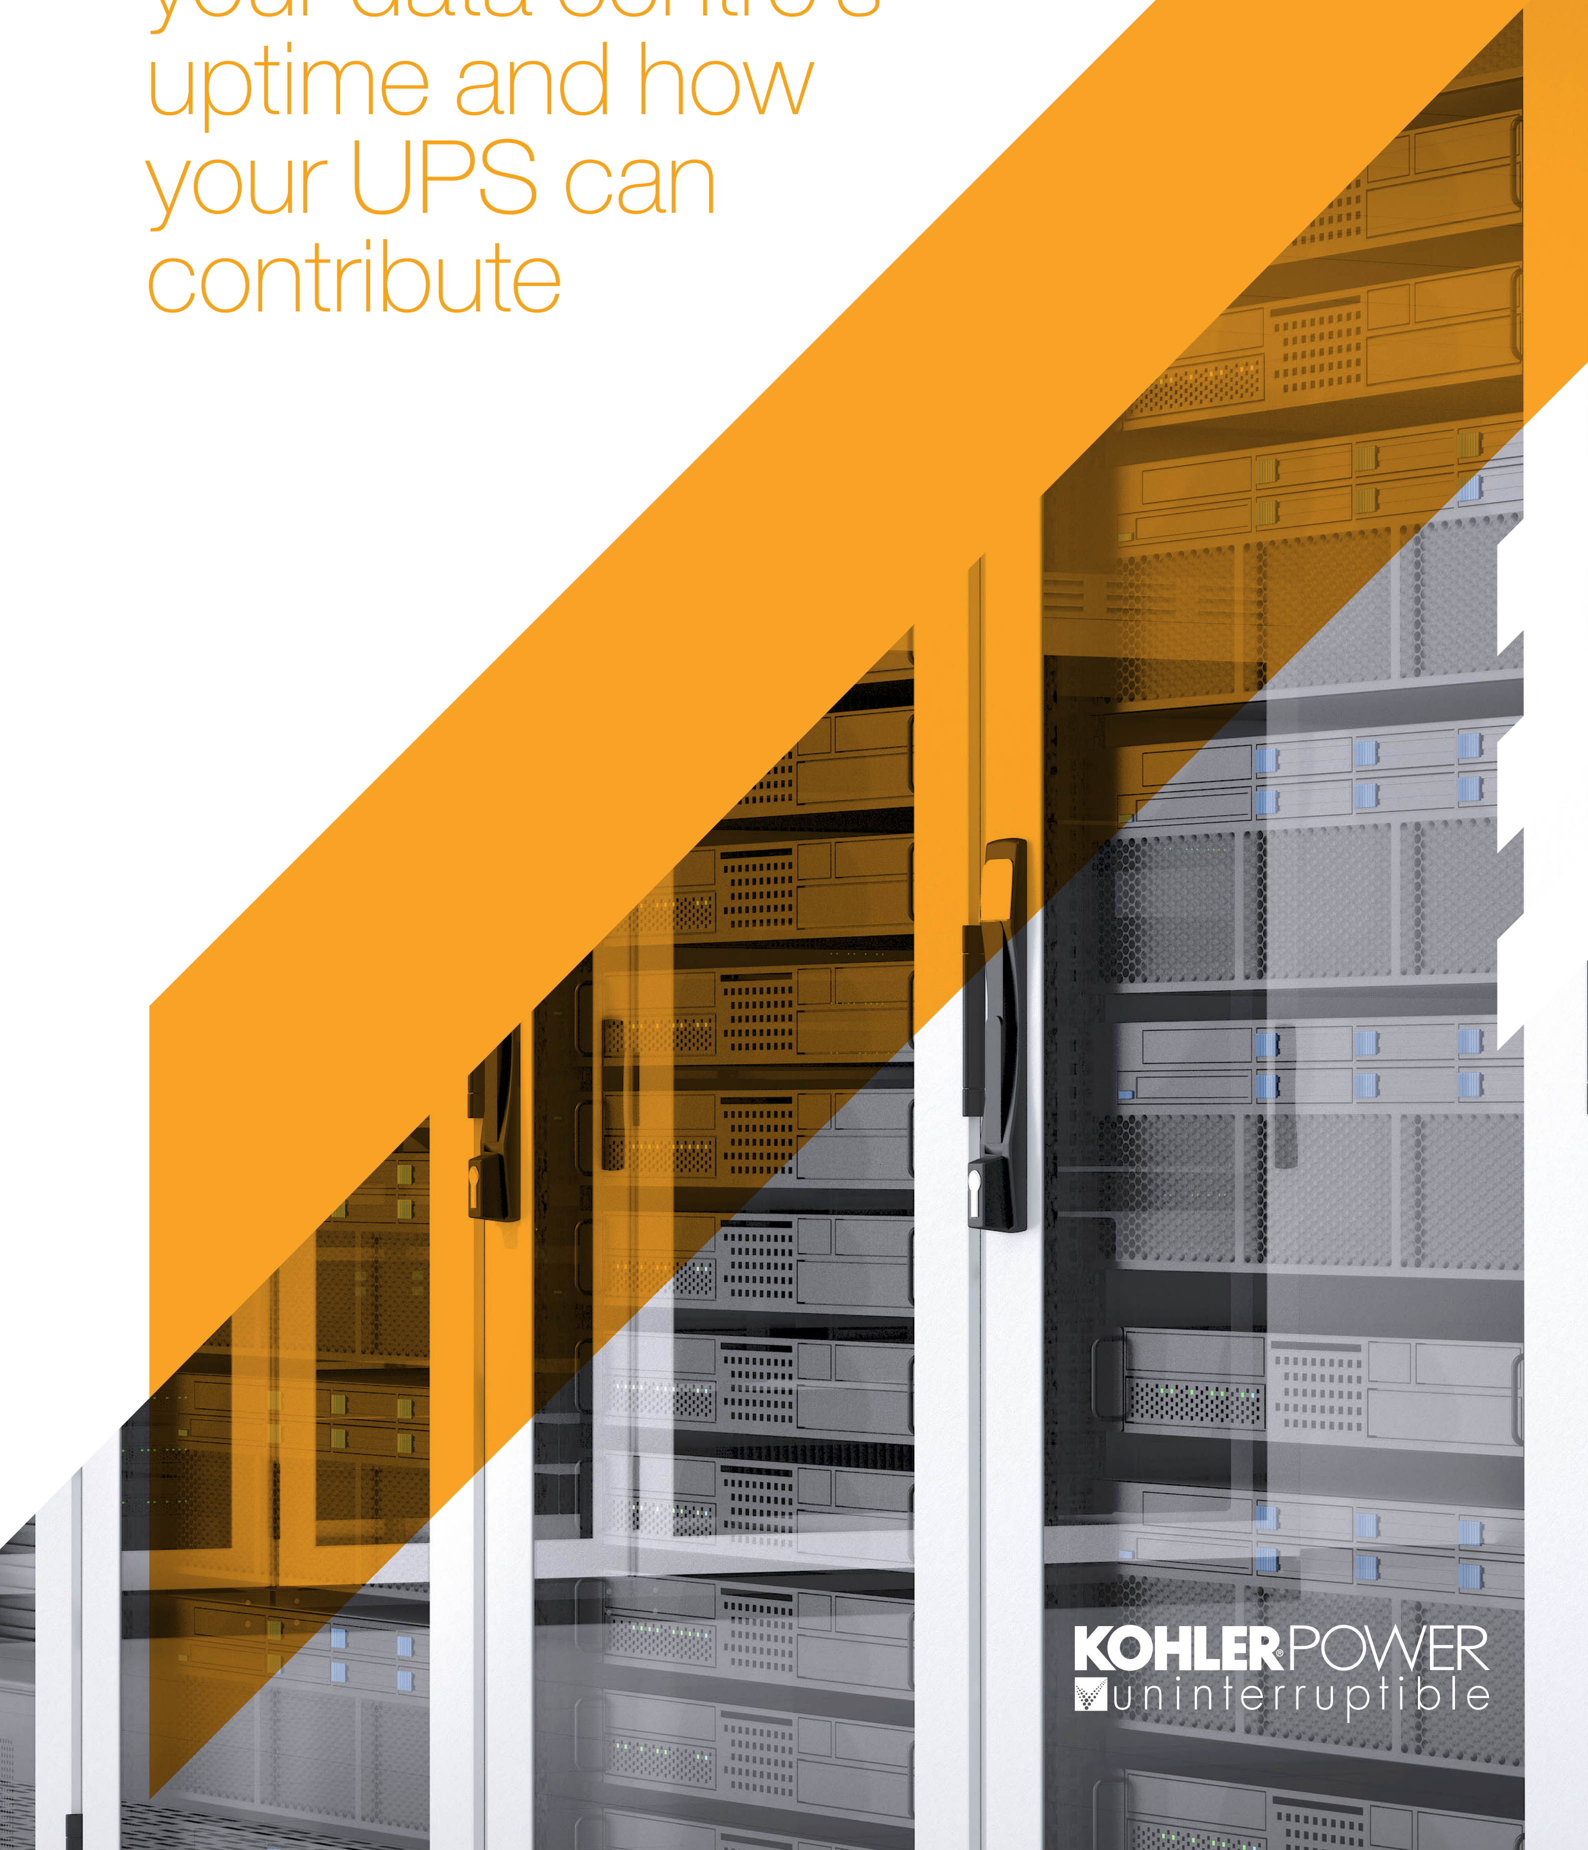

KOHLER POWER uninterruptible

Securing the highest availability money can buy is essential for on-line mission critical applications. However, there are also many less critical operations where a careful balance of sufficient availability against cost makes more sense.

Here, we look at the Uptime Institute's Tier Classification system, and show how this can be used as a tool to specify an appropriate level of data centre availability. This discussion includes UPSs and how they contribute to tier classification, both through internal topology and deployment.

Introduction

Is it the data centre operator's role to provide the highest ICT availability that money can buy, or does their primary responsibility lie in maximising their organisation's profit?

Not surprisingly, there's no single answer; instead, any data centre strategy will depend on its particular circumstances. In one scenario, for example, the facility may be being used to run a truly critical service, like an online banking operation with a fleet of ATMs. Here, no expense will be spared in securing the highest possible level of ICT resource availability. After all, the cost of any investment into maximised uptime will be dwarfed by the price of any failure, not to mention other consequential damage.

However, there are countless other scenarios, many of which are less critical. If data is being stored, say, for an offline bricks and mortar business, then 24/7 instant access may not always be essential – the business could ride through a loss of central ICT resource, at least for a limited period. In such a case, it wouldn't make commercial sense to match the previous example's level of availability; the increased infrastructure costs and operational complexity needed for extreme availability levels wouldn't be justified without the same risk of loss.

So, when data centre managers are planning and deploying equipment like UPSs, the most relevant question isn't simply 'What's the highest level of availability I can achieve?', but rather 'What level of availability is most appropriate to my business operation, and how do I achieve this?'

The Uptime Institute has been working on this question for a couple of decades – not just for UPSs, but for entire data centre infrastructures. However, UPSs do play an essential role within their considerations, as will become evident.

Overall, the Uptime Institute's experience is made available to data centre stakeholders in the form of a Tier Classification system. However, tier classification isn't a prescription for an annual availability figure; instead, it's about evaluating data centre performance in terms of uptime. It's a benchmarking tool that can be used by operators and users alike to assess a data centre's suitability for their business application and its criticality. This performance-based evaluation of a data centre's specific infrastructure is based on both design and operational elements.

This article is intended to give readers an introduction into how tier classification can help them evaluate their (or a third party's) data centre's suitability for their application, in terms of uptime performance. Then, as UPSs are one of the key factors that influence uptime and classification, we look at how they can contribute to facility performance at different tier levels through both their topology and their deployment.

Uptime Institute tier classification

Tier classification breaks down a data centre's potential for uptime performance into four levels, known as Tiers I – IV. The tiers are progressive, in that each tier incorporates the requirements of all lower tiers.

Tier I – basic capacity: A Tier I data centre provides dedicated site infrastructure to support information technology beyond an office setting. Tier I infrastructure includes a dedicated space for IT systems; an uninterruptible power supply (UPS) to filter power spikes, sags, and momentary outages; dedicated cooling equipment that isn't shut down at the end of normal office hours; and an engine generator to protect IT functions from extended power outages. It does not specify a redundancy concept.

Tier II – redundant capacity components: Tier II facilities include redundant critical power and cooling components to provide select maintenance opportunities and an increased margin of safety against IT process disruptions that would result from site infrastructure equipment failures. The redundant components include power and cooling equipment such as UPS modules, chillers or pumps, and engine generators.

Tier III – concurrently maintainable: A Tier III data centre must support equipment replacement and maintenance without need for shutdowns. A redundant delivery path for power and cooling is added to the redundant critical components of Tier II so that each and every component needed to support the IT processing environment can be shut down and maintained on a planned basis without impact on the IT operation. For Tier III, only one of the power paths has to be active at any time.

A Tier III site is susceptible to disruption from unplanned activities.

Tier IV – fault tolerance: Tier IV site infrastructure builds on Tier III, adding the concept of fault tolerance to the site infrastructure topology. It has redundant capacity systems and multiple active power distribution paths. Fault tolerance means that when individual equipment failures or distribution path interruptions occur, the effects of the events are stopped short of the IT operations.

To establish fault tolerance and concurrent maintainability of the critical power system between the UPS and the computer equipment, Tier IV ICT hardware must have dual power inputs.

The Uptime Institute comments that data centre infrastructure costs and operational complexities increase with tier level, and it is up to the data centre owner to determine the tier level that fits his or her business's need. A Tier IV solution is not 'better' than a Tier II solution. The data centre infrastructure needs to match the business application, otherwise companies can overinvest or alternatively take on too much risk.

Uptime Institute recognises that many data centre designs are custom endeavours, with complex design elements and multiple technology choices. As such, the Tier Classification System does not prescribe specific technology or design criteria beyond those stated above. It is up to the data centre owner to meet those criteria in a method that fits his or her infrastructure goals.

Uptime Institute removed reference to 'expected downtime per year' from the tier standard in 2009. The current tier standard does not assign availability predictions to tier levels. This change was due to a maturation of the industry, and understanding that operations behaviours can have a larger impact on site availability than the physical infrastructure.

Data centres can prove their status through a tier certification process. This comprises two parts: a Tier Certification of Design Documents (TCDD) and a Tier Certification of Constructed Facility (TCCF).

The TCDD involves a 100% review of the design documents, ensuring that each electrical, mechanical, monitoring and automation subsystem meets all fundamental concepts, and there are no weak links in the chain. TCDD is awarded when the Institute is satisfied that all the documents are entirely compliant.

However, it's also important to ensure that the site is actually built and operated as specified in the documents, and performs accordingly. This is determined by a site visit in which any discrepancies between the design drawings and the installed equipment are identified. When the data centre owner rectifies any such deficiencies, the Institute awards a TCCF certificate to complete the tier classification.

Beyond this, data centre operations can also be certified according to an Uptime Institute document, *Tier Standard: Operational Sustainability*. This is an important point, as it highlights the Institute's view that performance at any tier level depends on effective operations management as well as on infrastructure design.

So how does the tier classification system relate to users' business models? The Institute comments that Tier I and Tier II are tactical solutions, usually driven by first-cost and time-to-market more than life-cycle cost and performance (uptime) requirements. Organisations selecting Tier I and Tier II solutions typically do not depend on real-time delivery of products or services for a significant part of their revenue stream. Generally, these organisations are contractually protected from damages stemming from lack of system availability.

Rigorous uptime requirements and long-term viability are usually the reason for selecting strategic solutions found in the Tier III and Tier IV site infrastructure. Tier III and Tier IV site infrastructure solutions also have an effective life beyond the current IT requirement and are typically utilised by organisations that know the cost of a disruption – in terms of actual dollars – and the impact on market share and continued mission imperatives.

Figure 1: Uptime Institute tier certificate

Tier classification and UPSs

You will note that the four top-level tier definitions refer to UPSs in a few distinct ways:

Tier I infrastructures do not specify redundancy for availability improvement, but the definition focuses on the other aspect of every UPS's role; to protect the critical load from spikes, sags, momentary interruptions and any other power problems arising when the utility mains is available. Remember that this requirement propagates upwards through all the tiers, even if it isn't mentioned specifically.

Tier II introduces the concept of redundant components, and recommends that critical power and cooling equipment, including the UPSs, should be implemented as redundant modules.

Tier III builds on Tier II, and extends to redundancy in power paths as well in components. This does not necessarily affect an individual UPS's topology, but it will affect how a set of UPSs are deployed across a power network.

Tier IV has two active independent power paths simultaneously serving the ICT equipment, and 2N power and cooling systems.

UPS topology and availability

From the above, it's possible to visualise availability and redundancy at two levels. The first is within the topology of each UPS, so that any such UPS can withstand the failure of one or possibly more of its constituent modules. The second, at a higher level, applies to a redundant arrangement of multiple UPSs in a power network.

Accordingly, we can start by looking within the UPS, and seeing how its topology can be designed to maximise availability, and the role of redundancy in this. In doing so, it's essential to understand that availability is about the relationship between Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR). Specifically, the equation is:

$$A = \frac{MTBF}{(MTBF+MTTR)}$$

Where A = Availability.

This shows that we can increase availability either by increasing MTBF or by decreasing MTTR. An obvious way to increase MTBF is to improve the reliability of the UPS's internal components and design. While improving reliability does in practice increase MTBF, it does so only to a certain extent. As Figure 2 shows, the investment/MTBF curve starts to flatten out as spending on component quality is increased. Investment alone is often not sufficient to secure the high MTBF level required; MTBFs of 50,000 to 200,000 hours can typically be achieved through top component quality.

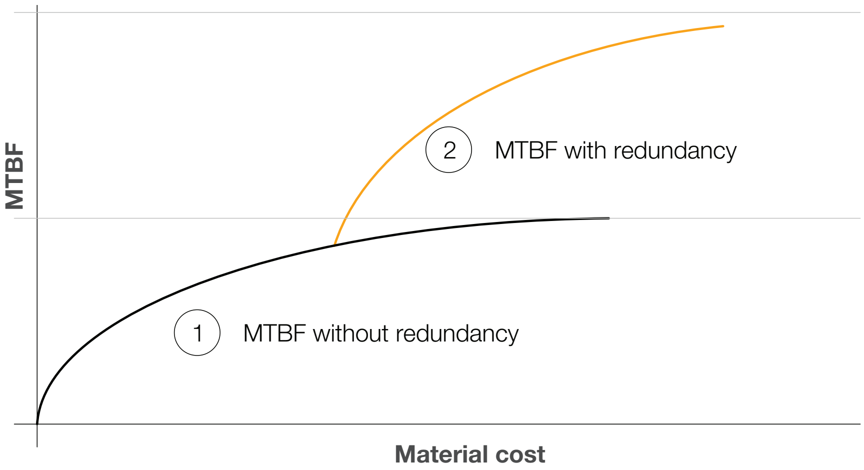

Figure 2: The effects of component quality and redundancy on MTBF

However, the graph also shows a second curve, labelled 'MTBF with redundancy'. This introduces the concept that you can also increase a UPS's MTBF by building in redundant components, because this allows resilience to one, or possibly more component failures.

Consider a simple example of a monolithic UPS system that has to support a 400kW critical load. The system could simply have a rating of 400kW, and it would support the load perfectly well – unless it suffered a failure, in which case the load would lose power immediately. Suppose instead that the same 400kW load is shared equally between two 400kW UPS units; if one fails, the other can continue to fully support the load until the fault is repaired. Redundancy has made the UPS resilient to a single internal failure.

Although this arrangement is attractive in that it significantly improves MTTR, it immediately doubles the UPS's capacity, heavily increasing both cost and footprint. Also, as each UPS unit never runs at more than 50% loading during normal operation, efficiency may be adversely affected. This suggests that a better approach might be to power the load from smaller UPS units in parallel, say, four rated at 100kW each. These will support our 400kW load, while incrementing to five will provide the redundancy element; in this case, N+1 redundancy, as the system has one module more than is required to fully support the load. This has been achieved with only 100kW instead of 400kW extra capacity.

Example 1	(1+1) redundant configuration – free-standing UPS units	(4+1) parallel redundant configuration – free-standing UPS units
MTBF	1,250,000h	500,000h
MTTR	6h	6h
Availability	99.9995%	99.9988%

Figure 3: Redundancy and availability figures for free-standing UPS systems

Example 2	(1+1) redundant configuration – free-standing units	(4+1) parallel redundant configuration – rack-mounted UPS modules
MTBF	1,250,000h	500,000h
MTTR	6h	0.5h
Availability	99.9995%	99.9988%

Figure 4: Impact of hot-swap modules on UPS MTTR and availability

Yet this solution, as represented so far, has a problem, as revealed by the table in Figure 3. This shows that the MTBF for the (4+1) configuration is considerably lower than for (1+1), because five modules have a greater potential for failure than two. Accordingly, its availability is lower than for the (1+1) solution.

However, we saw above how availability depends on MTTR as well as MTBF – and herein lies the solution. Modern UPSs are typically implemented with modular topology rather than as free-standing units. KUP's PowerWAVE 9500DPA UPS system, for example, comprises a frame that can accept from one to five 100kW modules. This offers many benefits in terms of efficiency, scalability and redundancy – and additionally, hot-swappability for its modules. A failed module can simply be pulled out of the 9500DPA frame, and replaced with a new plug-in unit, all without powering down or interrupting the load's supply. As a result, MTTR drops from the free-standing units' six hours to just half an hour.

The table in Figure 4 shows the impact of this on the UPS's availability. From this, we can see that a modular UPS system with hot-swappable modules can achieve an availability of 99.9999%, sometimes referred to as 'six nines' availability.

UPS Systems topologies – '2(N+1)' system

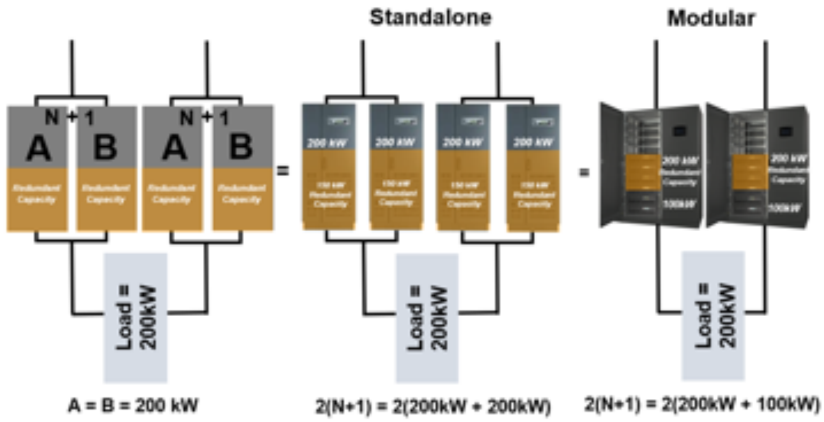

Figure 5: Redundancy at both UPS and power distribution path level

Fully redundant Tier III and Tier IV solutions

We have seen that Tier II upwards calls for redundant power components, while Tiers III and IV add a requirement for redundant power paths. Figure 5 shows an implementation that could satisfy a Tier III or IV installation, as it provides redundancy at both the UPS and power distribution path level.

Either UPS will survive a failure of a single 100kW power module, while the overall power system can tolerate the loss of either UPS without compromising power delivery to the load.

This 2(N+1) system provides concurrent maintainability with high availability; its benefits must be balanced against an increase in system complexity, and increased costs if not well utilised.

Conclusion

In this article we have shown the reasons for developing and using data centres that best balance cost against a particular application's risk, rather than simply investing in the highest possible availability. We have also introduced the Uptime Institute's four-level tier classification system, and its use as a tool by data centre owners and users to evaluate a facility's uptime performance and position within the cost/availability trade-off scale.

One aspect of tier classification is its description of how UPSs can contribute to facility uptime, both by their internal topology and by their deployment across the power distribution network. Accordingly, we have covered how these UPS aspects can help with achieving a data centre's required tier level.

This discussion has also revealed how the modular topology used in most modern UPSs simplifies the task of complying with Tiers II–IV redundancy requirements.

KOHLER POWER uninterruptible

Woodgate, Bartley Wood Business Park, Hook, Hampshire RG27 9XA

Tel: 01256 386700

Email: uksales.ups@kohler.com

www.kohler-ups.co.uk