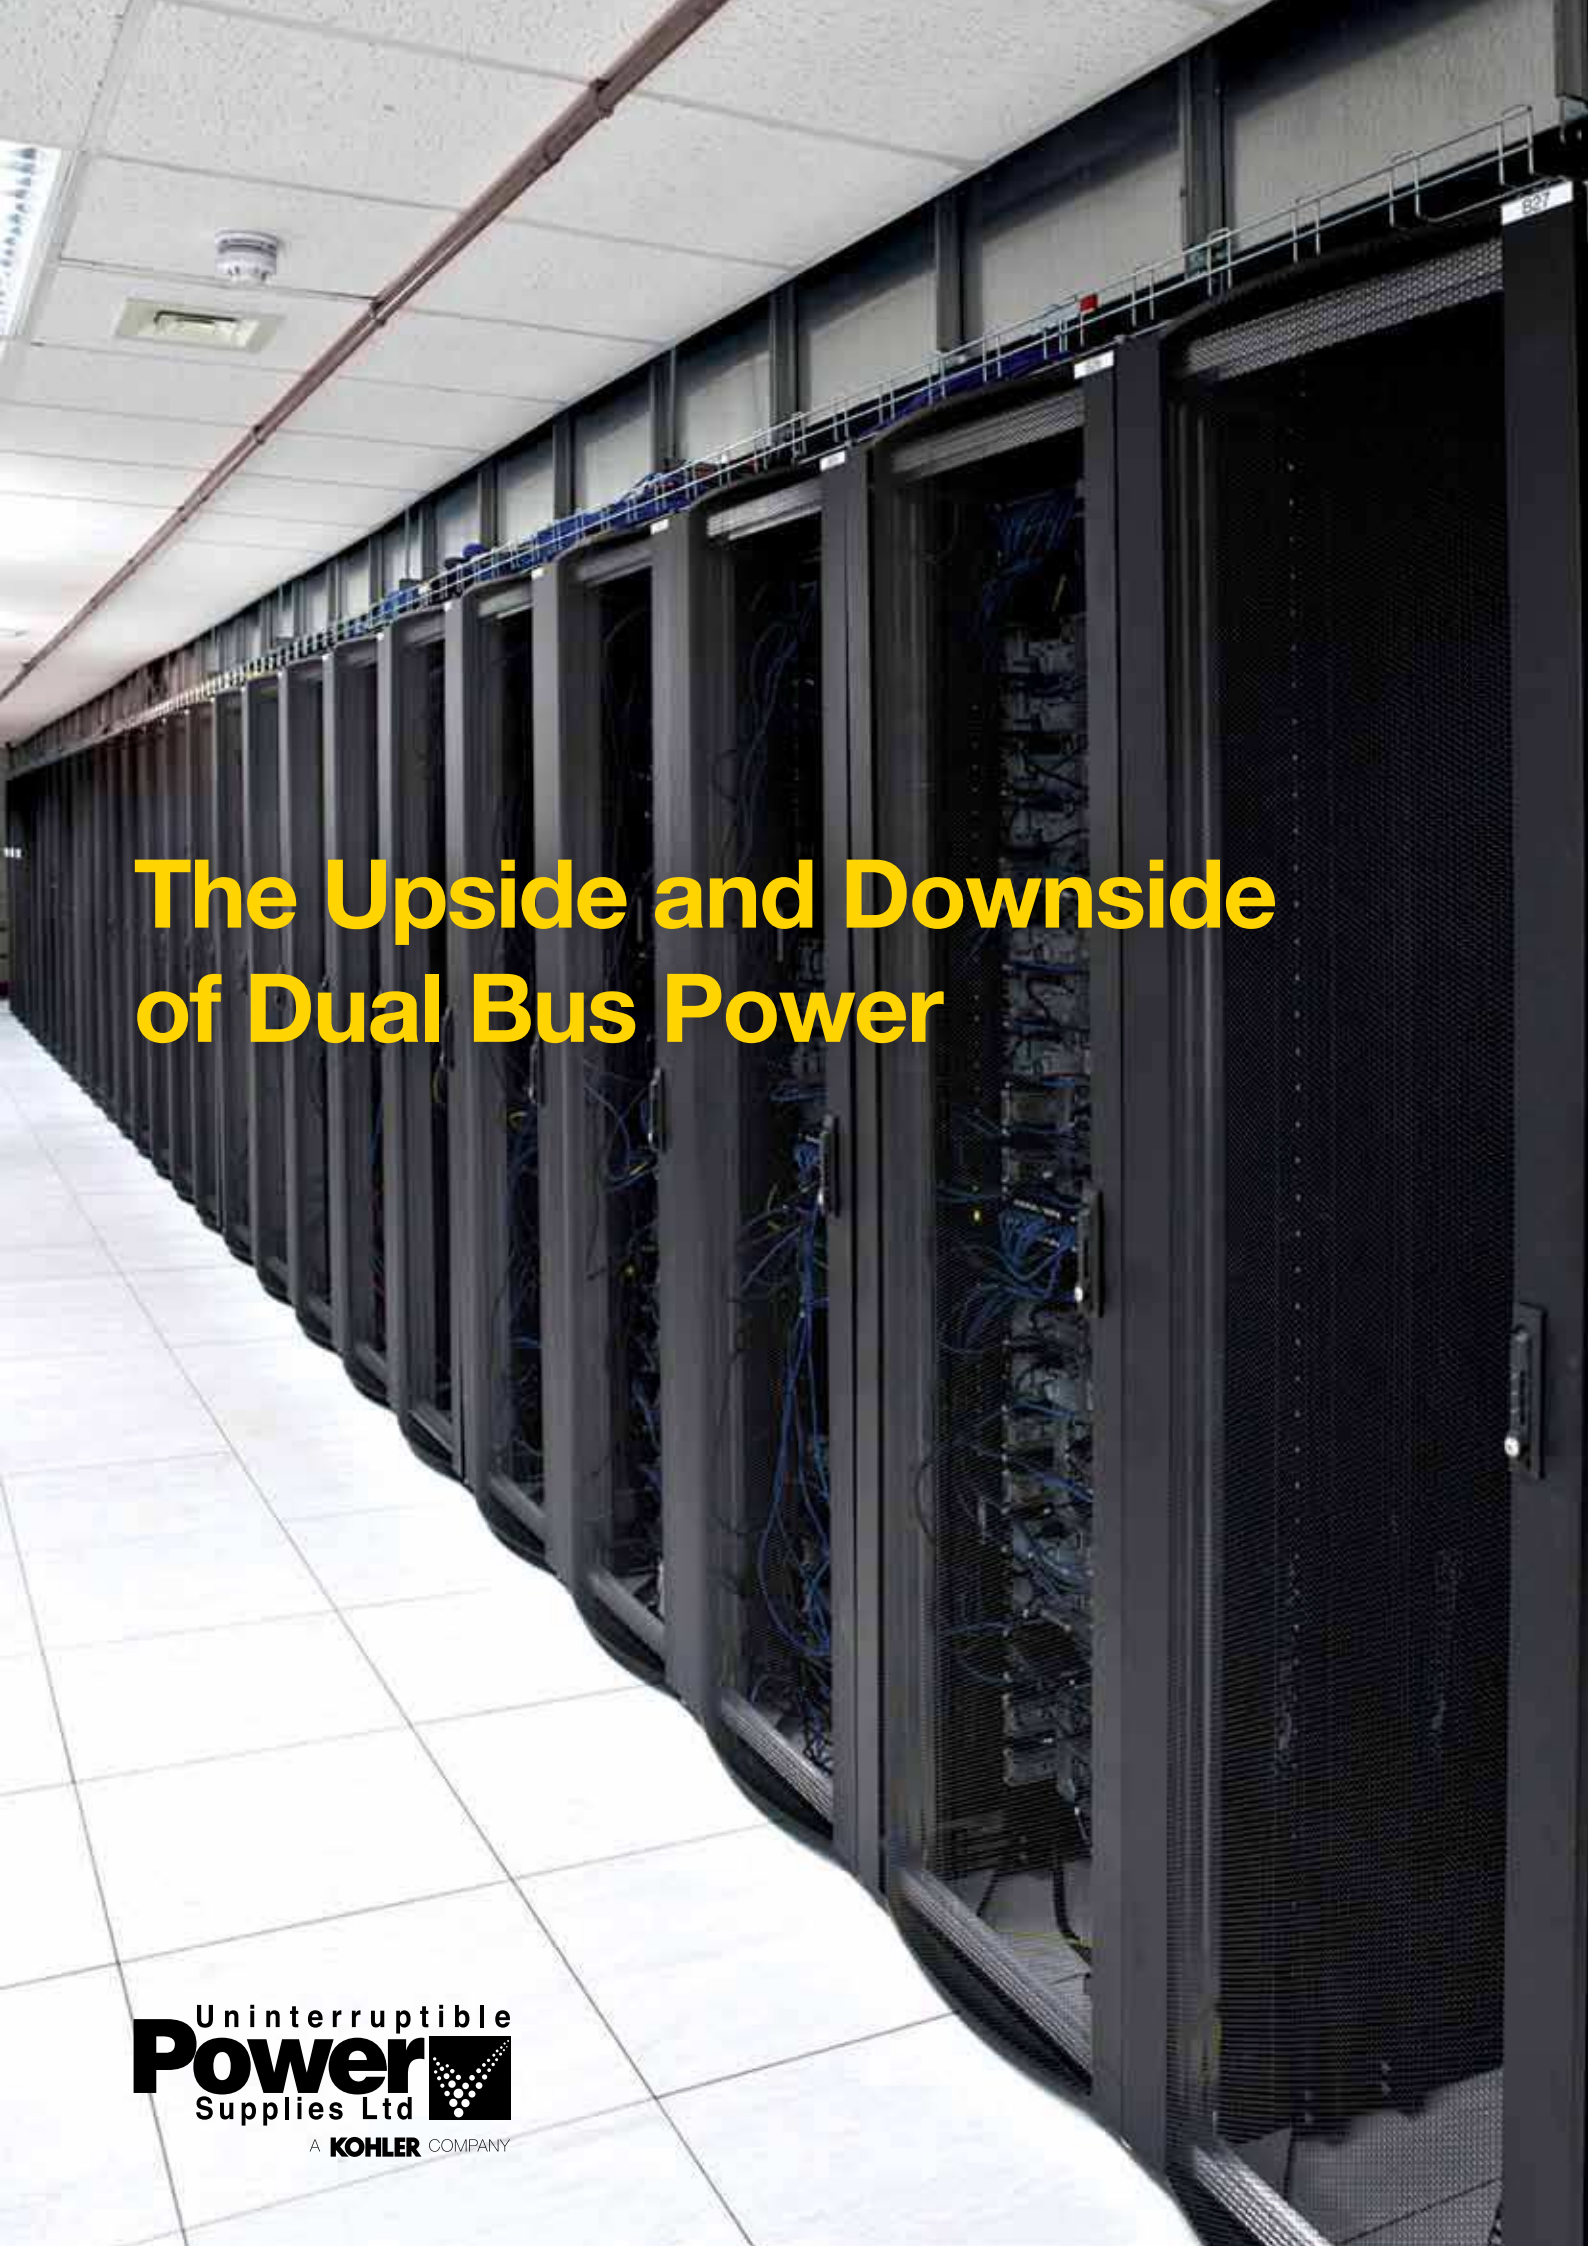

The Upside and Downside of Dual Bus Power

Uninterruptible Power Supplies Ltd

A KOHLER COMPANY

Introduction

Data centre facilities provide high-fidelity power to the critical load by the provision of Uninterruptible Power Supply systems in various levels of redundant architectures that are well described in the foundation work of The Uptime Institute in the USA. When the founders of The Uptime Institute introduced their data centre Tier Classifications in the early 90s they built on their own innovation of dual-corded ICT loads. Prior to that time ICT loads, such as enterprise servers, were single-corded devices and there were only two possible levels of provision in the power domain – firstly without redundant components, which often needed a load shutdown to carry out maintenance and where a single failure resulted in downtime, and the second with redundant elements which gave some opportunities for concurrent maintenance and a degree of fault tolerance. The best example is the UPS system – Tier I having a single module and Tier II having a redundant system with N+1 architecture.

With the advent of the dual-corded loads, the opportunity for concurrent maintenance expanded when Tier III introduced the principle of an active path (containing an N+1 system) to one cord and a separate passive path that brought power (from the utility or generator if required) to the other cord. For the ultimate reliability and resilience Tier IV brought active/active to the two load cords – with an N+1 redundant power system in each path, 2(N+1), that provided both concurrent maintenance and fault tolerance to the single major failure event.

It is interesting to note that once you assume that the most basic power system (Tier I) comprises a single UPS and single generator, there are only four possible power architectures to support a dual corded loads. Therefore it should not be surprising to see that many 'design' authorities followed in the steps of TUI and perpetuated the four Tier levels, e.g. TIA942, BICSI and the soon to be released EN50600. In the 20 years since the original Tier Classifications were innovated only one change has been seen – the reduction from 2(N+1) to 2N in Tier IV, although this has only been well described by the originators, TUI, and not followed by such standards as TIA942.

So what are the upsides and downsides of dual bus?

Upsides of dual bus topology

To move from Tier I to Tier IV clearly increases the potential availability in terms of both planned outage for maintenance and unplanned outage by failure, and the major step occurs between Tier II and Tier III as dual-corded loads offer the opportunity to utilise a dual bus power architecture. However the term 'Availability' is often misunderstood, usually misused and, sometimes, abused deliberately for marketing purposes. At the heart of this 'problem' is the original percentages that TUI published in their original white paper: For each Tier they gave an availability percentage and expressed it as 'X minutes downtime per year', e.g. Tier I offered 99.67% with 28.8 hours/year downtime compared to Tier IV, 99.99% with 53 minutes/year downtime. It should be obvious to all that one failure per year would be unacceptable for any system (I or IV) and the amount of downtime hardly matters when it may take the average integrated ICT load several hours to re-engage with the mission critical function after a loss in power.

Clearly the benefit of moving up the Tier layers is to extend the Mean Time Between Failure (MTBF) although if single-corded loads exist in the dual bus architecture then they should be protected by point-of-use (usually rack-mounted) static-transfer-switches.

After that (single) failure the recovery time (Mean Down Time) needs to be short as possible as, interestingly, to give an availability figure you need to know both the MTBF and the MDT, as follows;

$$Availability = \frac{MTBF}{MTBF + MTD} \; x100\%$$

It is worthy of note that any system requiring an annual shutdown of 4h for maintenance can only achieve an availability of 99.95% (MTBF=8,760h and MDT=4h).

So a high availability can be achieved by either a long MTBF or a short MDT but the MDT should (but usually does not) include the ICT system re-start time.

Having pointed out the weakness in the term 'availability' and accepting that MDT will always be several hours, we can better express the upside of climbing up the Tier layers as a relative MTBF of the alternative power system architecture. Figure 1 tabulates the relative MTBF of the architectures from N to 2(N+1) for a change in 'N'. In this case the MTBF describes the voltage supplied by the UPS system inside the latest version of the ITIC/CBEMA voltage tolerance curve, and ignores the MTBF of the downstream power distribution systems. In the case of the dual bus active/active the MTBF represents the event of concurrent failure of both buses.

So it can be seen that the MTBF of dual bus systems is dramatically enhanced over the MTBF of a single module. We can see in the last three columns a typical high power data centre (where N=2) the availability based on one failure event with a Mean Down Time of 8 hours – a 4h response on site followed by a 4h repair or an 8h re-boot time after a momentary failure in voltage lasting longer than 20ms.

However there is an additional advantage of any dual bus system over a single bus system:

Depending upon which analyses of data centre failure you choose to read, you will learn that 35-70% of all data centre failures are down to human error and most of those take place in the electrical infrastructure. The advantage of dual bus, be that 2N or 2(N+1), is that simultaneous human errors (i.e. one human error in each system at the same moment) is virtually impossible. The obvious cause of downtime in single bus systems is inadvertent operation of the EPO and that just can't happen in two separate rooms. So, the chances of human error are substantially reduced in dual bus systems.

It is worth noting that many data centre designs that are not Tier IV per se, incorporate 2N power to enable ease of management and maintenance without shutdown or risk. These are often referred to as Tier III+, although TUI do not support, in any fashion, the concept of intermediate steps in the Tier Classification hierarchy.

Relative UPS MTBF, CapEx & Availability

Architecture	Where N =			Where N = 2		
	N=1	N=2	N=3	CapEx	MTBF (y)	Availability (%)
N	1	0.9	0.8	1	10	99.991112%
N+1	10	9	8	1.8	103	99.999111%
2N	800	700	600	2.3	7,991	99.999989%
2(N+1)	1000	900	800	3.6	10,274	99.999991%

Figure 1: Availability calculated with a single UPS module MTBF=100,000 & MDT=8h

Historical downsides of dual bus topology

To counteract the clear advantages of dual bus there have been penalties. Of course, if the business model of the organisation requiring the data centre is centred only on ultra-high availability and high fault tolerance, these 'downsides' are the acceptable cost of doing business.

The most obvious penalty may be the initial capital investment in the extra redundant components, although the relative costs are outweighed by the huge increase in relative availability as indicated in the above table.

Additional plant-room space to house the transformers, generators, UPS, switchgear in segregated spaces and delivery paths, complete with environmental control, fire detection and suppression, lighting and security all add to the initial investment. Having said that, the facility costs represent less than 25% of the data centre 10-year Total Cost of Ownership but the cost of power, the next downside of dual bus architecture to be considered, is the largest element in the TCO – in some data centre business models as much as 50%.

If we consider the original 2(N+1) architecture at partial load, we will be able to see why the change to 2N came about as concerns about energy efficiency grew, even in the USA.

Figure 2 shows the configuration of 2(N+1) when N=2.

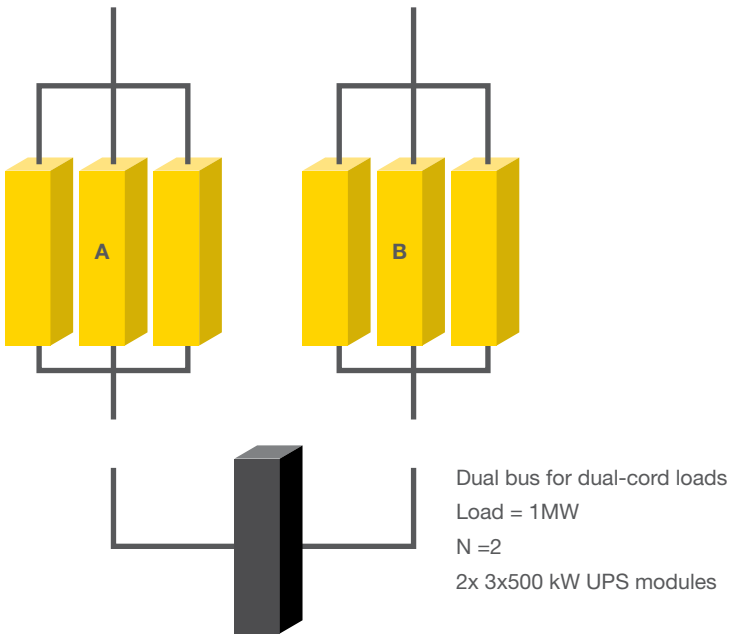

Figure 2: 2(N+1) when N=2

Let us now consider how an 'average' enterprise data centre facility load develops over time:

- The critical load often starts off below 20%
- It may climb to c40% after 2-3 years
- Reaches a plateaux of <80% after 5 years
- The critical load never reaches 100%

If we consider the 2(N+1) UPS system in Figure 2 the module load in normal operation against these indicative 20%/40%/80% load stages we can show that it is 3.3%/6.7%/13.3% respectively.

In other words, the UPS modules work at very light loads for the vast majority of their service life. Turning one module 'off' in each system to improve the load factor is only possible when the load is below 50%, and this improves the early years to 5% load per module at 20% facility load and 10% load per module at 40% facility load – still very low.

It was very probably this low load problem that pushed the change to 2N, removing the (double) redundant module on each side and relying on each system to be 100% redundant for the other. However, we need to view this against the efficiency of a typical North American legacy UPS: A typical large UPS in the USA was thyristor based, 460V input, 208V output with input, output and bypass transformers, often with a 6-pulse rectifier without a harmonic filter up to 600kVA. The efficiency at full load was c91% but the partial load efficiency was poor. On Figure 3 is shown (in red) a typical efficiency curve from c2005. To consider the impact of installing this type of UPS in 2(N+1) architecture we need to consider the efficiency in the sub-10% load range of around 50%. If we add a dual bus cooling system at similar low load it is easy to see that an operational PUE (rather than design PUE) of >3 was very easy to achieve.

Pressures on operational costs (power and maintenance) and the realisation that a double redundant system hardly gave any increased resilience for dual corded loads resulted in Tier IV being downgraded to 2N from 2(N+1).

Dual bus also can have an unfortunate side-effect outside of the UPS loading – that of under-utilised plant. This led to many forms of distributed redundancy architectures with 'swing' transformers or generators and often utilising static-transfer-switches. These solutions saved capital expenditure but at the risk of increasing complexity that sometimes led to lower reliability and the introduction of increased opportunities for human error.

One inadvertent consequence of following TUI and TIA942 recommendations for 15 minutes of battery autonomy per module in Tier IV systems, was that battery autonomy is not linear with load, and the effect of very light load on the dual bus system is to produce battery autonomies in the region of 3 hours. With power densities in the critical space gradually increasing, any extended UPS autonomy is not usable unless the cooling system is also continuous – since the load will shut-down on thermal-alarm before the power is shut-off.

Overcoming the problem – the modern solution for the legacy

There is no doubt that the extended MTBF and lower risk of human error that dual bus architecture offers with dual-corded loads, is as attractive now as it always was, if not more so. However it is now possible to mitigate, if not avoid altogether, the problems associated with partial load by utilising state-of-the-art UPS technology and designing the system in a scalable, modular, topology. Designing a data centre with dual bus power and N+1 cooling, both with concurrent maintenance capability, is generally referred to as Tier III+, despite TUI objections.

In Figure 3 we can see the dramatic improvement in efficiency, at all loads, between a legacy North American machine and an IGBT/IGBT transformer-free design.

In the same way that modern European UPS designs have dramatically improved in recent years, the change in part-load efficiency of static-UPS has also overtaken that of rotary-UPS in all its variants, so, for completeness, also included in Fig.3 is the curve for a typical hybrid-rotary UPS and a typical DRUPS.

The modern solution to this legacy problem is straightforward:

- Use high-efficiency IGBT/IGBT transformer-free UPS that has high partial load efficiency

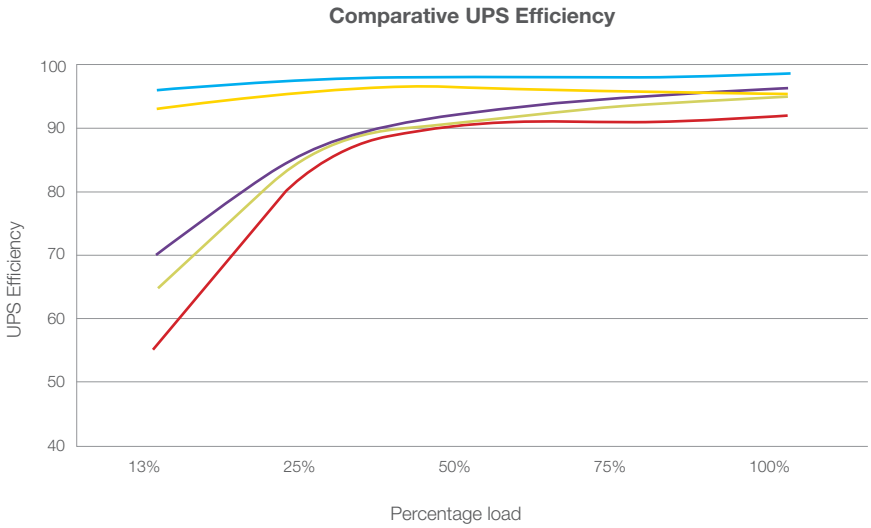

Figure 3: Comparative UPS Efficiency

- Apply it in a scalable way that suits the anticipated initial load and anticipated load growth profile – with the aim of maintaining as high a system load as is possible by turning 'off' any over-redundant capacity.

- In one (or even both) bus turn 'on' the UPS's eco-mode capability and virtually halve the system losses (ability to start the UPS inverter in static bypass if mode conditions dictate)

- Try to restrict the installed battery capacity on each module to under 10 minutes

For most large systems this strategy could result in an overall system efficiency of over 95% rather than 50% – a small price to pay for such high availability?

About Uninterruptible Power Supplies Limited

UPSL is a leading provider of power protection product and service solutions including UPS, standby diesel generators, battery banks, software and ancillaries. Part of the Kohler Corporation's Global Power Group, UPSL has a proven and enviable track record with customers across all market sectors around the world ensuring 24/7 power protection of mission-critical business systems.

Uninterruptible Power Supplies Ltd, a Kohler company,
Woodgate,
Bartley Wood Business Park,
Hook, Hampshire RG27 9XA

Tel: 01256 386700

Fax: 01256 386701

Email: sales@upspower.co.uk

Uninterruptible Power Supplies Ltd
A KOHLER COMPANY

UPS713-01-00 The Upside and Downside of Dual Bus – White Paper