Concurrent maintenance in data-centres: Tier IV updated

Ian F Bitterlin

PhD BSc(Hons) DipDesInn MCIBSE MIET MBCS MBIFM
Sales & Marketing Director

Prism Power Ltd, Watford, UK

Abstract

Data-centre power designs to Tier IV of the Uptime Institute & TIA942 are founded on the ability to carry out concurrent maintenance and the possession of fault tolerance. Following pressure from the energy efficiency lobbyists the requirement for Tier IV has been downgraded from 2(N+1) to 2(N). This paper reviews the original design, the criticism and the resultant changes.

A review of data-centre Tier Classifications

Information technology clients expect availability of 'Five-Nines', or 99.999%. Unfortunately, the substantial investment that a business makes to achieve Five-Nines in its computer hardware and software platforms is unlikely to be sufficient unless matched with a complementary site infrastructure that can support their availability goals.

The Uptime Institute[1] has, for more than 10 years, sponsored research and practical studies into data centre design, operation and resultant resilience and developed a Tier Classification to describe and differentiate facilities from an availability standpoint. A White Paper[2] from the Institute (authors of which include the originator of dual power supplies in IT equipment and the Tier system itself) is the basis of this review of the facility and operational concepts.

The overall site tier rating is dependent on all aspects of the site infrastructure and will be the lowest of the individual sub system ratings covering such aspects as power, cooling and distribution etc. It is important to be aware that sustainability (how the site is operated once constructed) also plays a significant role in what site availability is actually achieved. All too often people wrongly assume that installing a UPS is the end of their problems. However, if the overall design, installation and ongoing service support is handled badly, it could just be the beginning. For example, it is vital to ensure that the Mean Time to Repair (MTTR) the system is kept to a minimum if the highest overall availability is to be achieved. Nowhere is this more important than with the design of data-centers.

Each industry has a unique uptime-need driving the site infrastructure Tier level requirement: After careful alignment of IT availability objectives with site infrastructure performance expectations, an informed company may select a site infrastructure based on any of the tier classifications. Data centre owners have the responsibility to determine what tier of functionality is appropriate or required for their sites. As such, it is a business decision to determine the tier classification necessary to support site availability objectives. Part of this decision is to balance the IT operational practices with the facility practices that support the IT world. Once selected however, the desired tier should be uniformly implemented.
A more recent addition to the debate is a data-centre 'standard' in ANSI/TIA-942-2005 *Telecommunications Infrastructure Standard for Data Centers*, issued by Telecommunications Industry Association[3]. This follows the same Tier I-IV format and draws heavily on The Uptime Institute publications but extends the detail, especially in connectivity. It is entirely a USA centric ANSI specification, but can be used as a very useful

---

[1] The Uptime Institute, Building 100, 2904 Rodeo Park Drive East, Santa Fe, NM 87505, USA. www.uptimeinstitute.org
[2] *Industry Standard Tier Classifications Define Site Infrastructure Performance*; Turner, Seader & Brill, © 2001-2005 The Uptime Institute, Inc
[3] TIA Standards and Technology Department, 2500 Wilson Blvd, Arlington, VA 22201, USA. www.tiaonline.org

guide outside of the reach of ANSI. One point worth noting is that TIA-942 is specifically written for telecom related data-centre environments of a power density <2700W/m^2.

Getting to Five-Nines?

Maintenance and fault tolerance is the key to the Tiers and the table (right) shows the progressive level of redundancy and resilience required and how it might be achieved. This table refers to each of sixteen key systems that are identified by TUI as critical to the operation of a specific data-centre. For a facility to achieve a Tier classification it must achieve the benchmark in all sixteen criteria, of which Critical Power is just one.

	Tier I	Tier II	Tier III	Tier IV
Number of delivery paths	Only 1	Only 1	1 Active 1 Passive	2 Active
Redundancy	N	N+1	N+1	S+S or 2(N+1)
Compartmentalisation	No	No	No	Yes
Concurrently Maintainable	No	No	Yes	Yes
Fault Tolerant to Worst Event	None	None	None	Yes

Availability – a measure of 'goodness'?

To achieve a high-percentage Availability is simple; achieve a very long MTBF (Mean Time Between Failure) and a very short MTTR (Mean Time to Repair), the calculation simply being MTBF/(MTBF+MTTR) x100%.

TUI has assigned a target Availability (A%) to each of the Tiers and, sensibly, recommend to measure the downtime (MDT) over at least a five-year period.

It will be immediately apparent to the reader that to achieve a defined overall site Availability then each of the sixteen sub-systems must achieve much higher performance (e.g. A% raised to the power of sixteen). For the ultimate Tier IV site this means that every sub-system (e.g. power at the load terminals) has to achieve 99.9994% - the magic Five-Nines.

TIER	Site A%	Nines	MDT h/5y
I	99.670%	2	144.54
II	99.750%	2	109.50
III	99.980%	3	8.76
IV	99.990%	4	4.38

The importance of a short MTTR

Clearly a wide range of 'answers' can be generated by combinations of MTBF and MTTR (see right) but the reality is that only an emergency service back-up that can minimize travel time to site, have good parts availability and excel in first-time fix rate will achieve the sort of MTTR needed to push the Availability to the required level. Indeed it is quite easy to demonstrate that both Tier III and Tier IV cannot tolerate travel times of even 4 hours to site if they are to achieve the desired Availability performance – even with an MTBF in the 200-400,000h range.

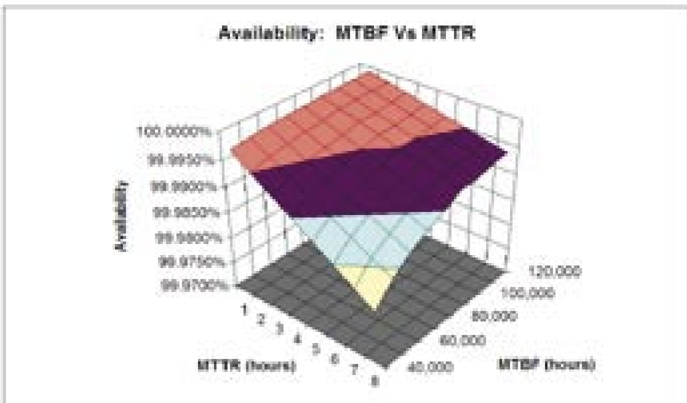

This conclusion raises the need for 24x7 remote monitoring, diagnostics and tele-assisted service via Internet enabled data-connectivity.

Higher Tier power systems don't come cheap…

Comparing systems is rather complicated when taking into account the type of load (single-corded or dual-corded) and whether or not Static Transfer Switches (STS) are deployed in the power system. In the strictest definition Tier IV is only intended for dual-cord loads without STSs, whilst they are essential in Tier III to transfer loads for maintenance routines. However at the most fundamental level we can take Tier I as the base cost and MTBF (=1) and make comparisons:

Tier	Load Cording	Concurrent Maintenance	A% Power System	MTBF Relative	COST Relative
I	Single	No	99.98333%	1	1
II	Single	Very limited	99.98547%	1	1.49
II	Dual+STS	Limited	99.99965%	45	1.65
III	Dual+STS	Partial	99.99983%	45	1.97
IV	Dual	Yes	99.99999%	2,450	3.11

Tier IV, for dual-corded loads without STSs is, by more than 50x, the most resilient power architecture possible. The drawbacks are the high CapEx (a 50% premium over Tier III), higher OpEx with partial load inefficiencies and under-utilized plant that can be regarded as wasteful of resources.

If the client needs a specific classification (e.g. Tier IV for a given business case) then there is little choice but to follow the guidelines in TUI &/or TIA-942. However, if the client does not insist upon a specific Tier classification it can be shown that other schemes (non-Tier, all involving STSs and higher plant utilization) can offer 8-10x improvement on Tier III performance at only a 10-20% cost premium.

There is a strong reliability and availability case to be made for having 2xTier-I facilities (several km apart) instead of 1xTier-IV facility – as long as the IT strategy supports parallel computing.

So what was 'wrong' with Tier IV?

For dual-corded IT hardware there really is little choice if the user wants continuous (24x7xForever) computing than to have a dual-bus power system – regardless of the cost and space it consumes, although it is worth noting that there are (despite marketing claims to the contrary) very few 'real' Tier IV facilities in the world.

However in its original format, which required two active paths each with N+1 redundancy, it led to systems running at very low load and, consequently, very low efficiency.

As an example of how bad it can get in the power system consider a 500kVA design load fed by two systems each of which comprises 2x500kVA UPS modules in N+1 configuration. Tier IV dictates a maximum operating load of 90% - but let's consider the more usual case of the load being 40% after the first year of operation in a new centre. In normal operation each path carries half the load, 20%, and each UPS module will carry half of that, 10%. The average UPS will have a 10%-load efficiency figure of less than 70%. That is a huge waste of energy – and the same applies, if not more so, in the mechanical cooling systems.

Pressure for change and the response

As 'green' issues climbed up the corporate agenda (even in the USA) the apparently wasteful over-capacity (through over redundancy?) of the upper Tier classifications came under the spotlight and consultants worldwide found it increasingly difficult to meet the needs of their clients when energy efficiency was on the wish-list. The Uptime Institute, always at the forefront of the data-centre infrastructure industry responded with an inspired piece of analysis.

Full membership of TUI is reserved for end-user organisations that own &/or run data-centers. Part of that membership is the confidential individual contribution, consolidation, analysis and re-distribution of performance metrics for the facilities. The collective data is for the benefit of the members and, apart from the occasional snippet that finds its way out into the wider world, is managed on their behalf by TUI.

It was therefore a relatively simple exercise (albeit time consuming) for TUI to compare members 'fault' reports with their power system configuration. The result was clear and published in a revised White Paper: There is no appreciable difference in IT-function failure rate between those systems that have 2(N+1) and those that have 2N.

In our example (previous section) that would double the load per module and raise the operating efficiency by at least 10 percentage points.

Is this a surprising finding?

It certainly impacts upon TIA-942 since that standard drew heavily (if not entirely) on the original TUI White Paper that specified 2(N+1) for Tier IV but, if looked at from a distance, the conclusion should not be surprising:

There is, obviously, going to be a difference in the theoretical MTBF between two power paths of 2N and 2(N+1) topology. A simple calculation - based on a single module MTBF of 80,000h, an N+1 system having an MTBF of 10x that and 0.5x for a power-parallel N solution can show that the potential improvement is dramatic – 400x.

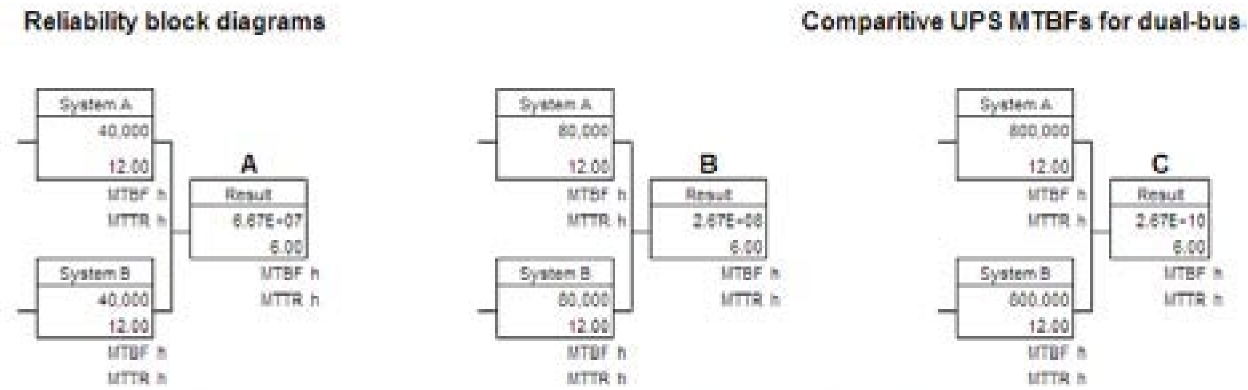

System	Description	MTBF		Relative MTBF
A	2N where N = 2 UPS in power parallel	6.67E+07	h	1
B	2N where N = 1 UPS module	2.67E+08	h	4
C	2(N+1) where N = 1 UPS module	2.67E+10	h	400

All calculation techniques in accordance with BS5760 Part 9:1992

So why is the reality so different? The answer probably lies in the magnitude of the answer compared to the population age: Consider that a facility lasts 10-15 years, so the average age of the TUI membership population is going to be around 7 years but the minimum MTBF (system A above) is 7,614 years – so the potential for a 400x improvement is unlikely to show up for a while. Add to that the relatively small population of only a few thousand facilities and the lack of apparent advantage for 2(N+1) over 2N is not surprising at all.

Conclusion

The change of the requirements for Tier IV redundancy is welcomed by the green lobby and many engineering consultants alike. Clearly, The Uptime Institute has reacted in a responsible and timely fashion to the consequences of its original work in the light of an ever-improving energy-efficiency environment. The updating of TIA942 is now required.

End