

BACHMANN Whitepaper

Redundancy and monitoring at rack level

A safeguard against failures in data centres

Redundancy is often underestimated

When the word “redundancy” is uttered, many IT professionals will immediately think of N+1-redundant UPS and CRAC units, or of multiple power supplies in servers and switches. Yet it is not merely the presence of redundant devices or components that provides the solution. Instead, monitoring of threshold values is what ensures true redundancy. Particularly at rack level (i.e. for end devices), the issue of redundancy is all too often underestimated or even misinterpreted.

A question of power [kW]

The question of how much power to allocate to each server rack already crops up during the planning phase for data centres. Due to the rapid growth in the IT sector and the constant increase in IT applications, it is getting more and more difficult to make reliable forecasts for power requirements after 2-3 years, let alone for the next 10 years. One reason for these forecasts being so difficult is that the assessment of electrical power requirements depends on the server utilisation. Since the introduction of virtualisation and cloud computing in particular, it has been getting ever trickier to evaluate the data centre architecture as a whole.

When providing specifications on the type plates of their products, IT equipment manufacturers already include a large safety margin for the power consumption [kW] or current consumption [A]. These generous specifications are one reason why planned figures for each rack’s power requirements are often wrong.

Recommendations for planning

Consultation with several planning companies and IT system integrators alike gave rise to the following procedure: taking the specifications on the type plates, a figure of 75% is assumed as a benchmark for the actual maximum values.



Certain IT users then assume 75% of this value for the simultaneity factor. Using the Hewlett Packard type plate shown, this would therefore mean the following:


Specification on the type plate:	575 W (max.)
Assumption of 75% as the real value:	431 W
Simultaneity factor in the rack = 75%:	324 W

With 10 server units of this type in the server rack, 3,240 W of power would therefore be required.

Redundancy with A and B infeed

Typically, electricity is fed to the IT equipment in a server rack by means of an A/B supply. Modern IT equipment generally features two (although sometimes more) redundant power supplies. In other words, the required electrical power is distributed equally between both power supplies: in this example, that would mean around 1,600 W or 1.6 kW for each.

Redundancy is only guaranteed if one PDU is able to assume the capacity of the second PDU in the event of escalation without slipping into the threshold load range. Over the past few years, an empirical value of 40% has become established. With a total current requirement of 14 A in the server rack (for approximately 3.2 kW), this would mean a maximum load of 7 A per PDU. This is the only way to guarantee that the second redundant PDU will cover the capacity of the others in the event of failure in one current path (A or B).



Author: Tilo Püschel, © Bachmann GmbH 2020

Bachmann GmbH
 Ernsthaldenstrasse 33
 70565 Stuttgart
 Germany

Tel. +49 711 86602-0
 Fax +49 711 86602-34
 Email info@bachmann.com
<http://www.bachmann.com/>