DATA CENTER

Exploring Dual-Power Feeds in Data Centers

Denis Blouin | 05.17.2018

Related Tags

Power & Monitoring

Share

Things like always-on technology, streaming content and cloud adoption are creating high demand for efficient, resilient and fast data centers that never let us down.

To meet these needs, dual-power feeds – two independent electrical feeds coming into a data center from the utility company – are becoming more common to reduce the chance of a complete outage (or not having enough power). This type of power set-up is often seen in Tier 4 data centers. If one of the two power sources suffers from an interruption, the other source will still supply power.

Generally labeled "A" and "B" feeds, each power source has not only its own utility feed, but also:

- A backup generator
- A switch that alternates between A and B feeds
- Electrical and distribution switchboards
- An uninterruptible power supply (UPS)
- A power distribution unit (PDU)
- Rack-level PDUs

At any one of these points along the chain, failure can occur. A true dual-power feed means that there are two separate sets of these components operating independently, reducing the likelihood of downtime due to failure.

Today, most mission-critical IT equipment, such as servers and switches, are also designed with at least dual power supplies. When everything is running normally, the equipment pulls power equally from both power feeds. In the event of an outage, however, the IT equipment can automatically switch all power to one feed or the other.

Using PDUs in a Dual-Power-Feed Environment

To make working on power supplies and inside your racks easier, and to reduce the opportunity for human error, some data center managers choose to color code rack PDUs to clearly identify A and B feeds. This eliminates confusion and helps prevent outages that lead to downtime – and also tells data center managers exactly which PDUs are connected to which power feeds.

Although it's common to see two PDUs per rack in a data center, we've also seen situations where – due to power demand – there are four PDUs per rack (two for the A feed and two for the B feed).

Having two (or four) rack-level PDUs available for a server means that the computing load is shared – when two are used, each PDU handles about 50% of the load. When one power feed fails, the remaining supply supports the entire load. For this reason, despite what it may seem, a PDU becomes overloaded at any level above 50% power. If IT equipment experiences a power supply failure,

For this reason, Belden recommends not surpassing 50% of a rack PDU's rated power load. Metered PDUs can help with this by monitoring load levels and alerting you if they reach certain power levels.

If you have questions about PDUs, want to make sure yours are set up correctly or need to verify that your PDUs aren't surpassing their rated power loads, Belden can help. Learn more about our customized data center infrastructure solutions.

Previous Post
Data Center Infrastructure: The Importance of Accuracy

Next Post
Label Your Data Center and Network Now – Benefit Later

About the Author

Denis Blouin

Product Line Manager

Denis is a product line manager for Belden R&E folio. Denis holds a BSc in Mech Eng (1989), His focus at Belden is helping data center managers find solutions to density challenges (cable mng, heat, power). He has been involved with the deployment of over 3 million square feet of white space. In his spare time, Denis enjoy golfing and MTB.

Contact Us

Find an Installer

Find a Distributor

Subscription Center

About Belden +