



CONSIDERATIONS FOR A HIGHLY AVAILABLE INTELLIGENT RACK POWER DISTRIBUTION UNIT

A White Paper on Availability

Introduction

Data centers are currently undergoing a period of great change. Data center managers are struggling to keep pace with growing capacity needs while working under the constraints of tightened budgets and energy efficiency initiatives and the challenges presented by new technologies such as virtualization and cloud computing are transforming

As data center environments become more dynamic and complex, many organizations are taking a more proactive approach to management and gaining better control of their data center operations so they can maintain or improve availability in increasingly dense computing environments while reducing costs and boosting efficiency. One area of improvement is inside the rack with the growing importance of intelligent Rack Power Distribution Units (rack PDUs).

As the last link in the power chain delivering critical power to IT loads, intelligent rack PDUs are a strategic asset for achieving high availability through elevated levels of responsiveness to change in data center capacities and densities. The emergence of data center infrastructure management (DCIM) is further increasing the role of the intelligent rack PDU within the data center. Data center managers are taking advantage of the benefits provided by the technology, including access to rack-level and IT equipment power consumption, visibility into rack-level environmental conditions, the ability to directly control power to IT equipment and rack-level capacity and power management.

This white paper discusses the considerations that need to be made when investing in intelligent rack PDUs to ensure that they do indeed provide a high availability solution.

Five aspects of a high availability design for Intelligent Rack PDUs include:

1. Reliability
2. Functionality
3. Fault Tolerance
4. Maintainability

Reliability

Intelligent rack PDUs (Figure 1), offering flexible and comprehensive remote management and real-time monitoring capabilities, can provide the best view of IT power consumption and rack operating conditions. However, while these new rack PDUs offer advanced capabilities, they are not much use if their core function is compromised or impeded: providing basic power distribution in any circumstance. There are a number of factors that should be considered to help ensure the primary function of intelligent rack PDUs.



Figure 1: Today's intelligent and adaptive rack PDUs provide value beyond just power distribution enhancing business agility, efficiency and availability.

High Temperature Rating

Given their location at the back of the rack toward the hot-aisle, rack PDUs are exposed to some of the harshest temperatures found in the data center. (Figure 2) In this location, it is fairly common to see temperatures of at least 122 degrees Fahrenheit (50 degrees Celsius). As power densities continue to increase and more organizations consider increasing the temperature in the data center to reduce energy consumption costs, these temperatures are only expected to rise. It's important that intelligent rack PDUs are rated to withstand a temperature of 131 degrees Fahrenheit (55 degrees Celsius) or higher.



Figure 2: Located at the back of the rack toward the hot-aisle, rack PDUs can be exposed to temperatures exceeding 122 degrees Fahrenheit (50 degrees Celsius).

power consumption of the rack PDU significantly lower. A switched rack PDU offering with bistable relays can save a typical 100 rack datacenter customer up to an estimated \$4660 annually in energy consumption costs.

Appropriate Overcurrent Protection

For safety reasons, regulatory agencies require rack PDUs to have overcurrent protection (OCP) above 20 Amps. If the appropriate OCP is not used in the data center, it could lead to incidents of tripping that compromise the availability of all loads connected to a branch circuit. An appropriate OCP for rack PDUs should not be highly sensitive and should have a minimal mean time to repair (MTTR). There are different types of OCP devices that can be used with intelligent rack PDUs. These include fuses, thermal magnetic circuit breakers and hydraulic magnetic circuit breakers.

Due to the downtime involved in replacement of **fuses**, some rack PDU manufacturers recommend that they not be used for mission critical facilities such as data centers. If a fuse is blown, it must be replaced – which can be a very time consuming and expensive fix. In most cases, this involves the upstream breakers on the floor PDU being turned off, a step that typically requires a qualified electrician. The result is significant downtime and a longer MTTR.

Circuit breakers are better suited for higher density, higher power consumption applications, primarily because they can be quickly and easily reset. **Thermal magnetic breakers** are designed to trip instantly as soon as the current threshold has been reached. They are also more sensitive to ambient temperatures, which does create an issue given the location of the rack PDUs. **Hydraulic magnetic breakers** are more tolerant of current surges and less sensitive to ambient temperature changes, making them an ideal choice for intelligent rack PDUs.

Another important point to consider about OCP is the **branch rating**. Branch OCP devices found in most rack PDUs are rated for either 80 percent or 100 percent of their load. This means that for a 20A OCP rated at 80 percent, the maximum continuous current that it can be used for is only 16A. An OCP rated at 100 percent would provide a maximum continuous current of 20A, making it an ideal choice to minimize the chances of tripped branch circuit breakers due to minor overloads. As always, any circuit breakers you choose should carry the appropriate agency approvals, such as UL489 within North America.

Low Idle PDU Power Consumption

As organizations invest in higher density IT equipment, the number of components contained inside a rack PDU has increased in order to provide higher levels of intelligence. However, the overall form factor has remained relatively small. It is important to realize that higher power consumption of the rack PDU itself leads to higher internal stress. The amount of heat dissipation within the rack PDU is determined by its idle power consumption. If the idle power consumption level is lowered, the heat dissipation within the rack PDU is lowered. Idle power consumption should be an important consideration especially for Switched rack PDU designs, which provide the ability to turn on, turn off or recycle power to connected IT equipment through the use of relays at every outlet. Idle power consumption is generally not a specification put out by vendors and hence should be requested.

Below are key features that help keep the idle power consumption of intelligent rack PDUs low.

- Aluminum construction - Aluminum has better conductivity than steel and helps keep the internal temperature rise to a minimum. It also makes the rack PDU lighter in weight, resulting in easier deployments.
- Backlit LCD display
- Bistable relays - These relays, also known as latching relays, draw power only when there is a change of state outside of normal operation. This helps keep overall

Smart Inrush Current Management

Inrush currents are caused by bulk capacitors charging in server power supplies, and can be in excess of 50A for a few tens of microseconds. To ensure that any upstream breakers do not trip, switched rack power distribution units should be considered since they allow the capability to power up the outlets in a sequential fashion.

These high inrush currents can also be detrimental to the relays within the switched rack PDUs themselves. Key to smart inrush current management within switched rack power distribution units is to ensure that the opening and closing of the relays is synchronized to be near zero crossing of the current / voltage waveforms.

Input Cable Size

When choosing Wye-connected rack power distribution units, ensure that the input power cord is properly sized to handle neutral currents in the event of unbalanced loads. This is especially important to consider outside of North America, since most power distribution units sold in those regions tend to be Wye-connected.

Power Cord and Outlet Locking

Power cord and outlet locking mechanisms secure the physical connection and ensure the power cords are not accidentally pulled out of the outlet, causing an inadvertent load drop. (Figure 3) Globally, the most common standard for outlets used in rack PDUs is IEC320 C13 and C19. IEC receptacles are internationally acceptable and handle output voltages up to 250V.



Figure 3: Locking outlets and locking power cords prevent accidental unplugging of IT devices.

Functionality

Intelligent rack PDUs should be able to provide proactive notification of impending issues before they occur. Warning and critical threshold settings for the current ensure that the rack PDUs do not experience overload conditions that could otherwise trip the breaker and the connected loads. While setting the current configuration, care should be taken that in a typical 2N scenario at the rack level, the thresholds for the branches are set at less than 50 percent of the overall rack PDU rating.

Software Electronic OCP

Paired with proactive monitoring, this feature will turn off and lock down all unused outlets on a branch circuit that has exceeded current established thresholds. It basically prevents someone from plugging new equipment into an unused outlet and causing a circuit overload.

Additional parameters that an intelligent rack PDU needs to monitor to ensure high availability:

1. Phase currents, along with a notification of unbalanced loads.
2. Temperature within the rack, through integrated sensors, along with the ability to configure auto turn-off outlets when temperatures exceed critical thresholds.
3. Ability to monitor circuit breaker status. (This is typically found in rack PDUs with metering or switching capabilities down to the outlet level. For rack PDUs with metering capabilities only at the branch circuit level, low critical threshold could be monitored as a proxy for circuit breaker status.)

All notifications should be capable of being received in a familiar format, such as SMS, SNMP traps or e-mail. Power distribution units should be capable of integration with a centralized management software, which will enable them to be easily managed.

Fault Tolerance

Intelligent rack PDUs should be designed such that a loss of a single phase will not lead to dropped power on all the unaffected phases. Also, regardless of the advanced features offered by an intelligent rack PDU, it needs to continue to provide basic power distribution in the event of a compromise of the intelligence capabilities. Fault tolerance due to a loss of the one of the main intelligence capabilities (i.e. switching, metering and external connectivity) is based on the design of these capabilities.

Metering

Current sensing within electrical circuits can be provided through the use of shunts, current sensors or Hall Effect sensors. Because shunts sit in the path of high voltage power, an issue with the shunt itself often leads to disruption of power within the primary circuit. On the other hand, current transformers and Hall Effect sensors are coils that are isolated from the primary high voltage circuit. Therefore, a disruption of power to these sensors themselves has minimal impact on the power flowing through the primary path. Current transformers have a further advantage over Hall Effect sensors in that they provide higher accuracy.

Switching

Switching within rack PDUs is vital to being able to remotely turn on or off the connected equipment, and is enabled by the use of relays at each outlet. The relays used within rack PDUs can be one of three types: normally-open; normally-closed or bistable relays.

- Normally-open relays require power to be supplied to them in order for the outlets to be able to deliver power to the connected loads. In the event there is an issue with the power supply feeding the relays, the relays remain open and there is no power provided to the connected outlets.
- Normally-closed relays only require power to open the outlets. Under normal operation, they remain closed. This means if there are any issues with the power supply feeding the relays, the outlets continue to provide basic power distribution to the connected loads.
- Bistable relays (or latching relays) act as normally-closed relays during normal operation, in that they also resort to providing basic power distribution in the event of a power supply failure. They do provide an additional advantage in that they allow a choice of state when power is restored after an outage. Outlets can either be turned on or off, or returned to the same state they were in prior to the outage. Bistable relays only require power if they change a state. They will keep outlets up and running with no power. In addition, their power consumption during normal operation is low, which helps minimize the overall energy footprint of Switched rack PDUs. (Figure 4)

External Connectivity

In the event that the primary network to the rack PDU goes down, some rack PDUs provide redundant communications through integration with out of band management devices, such as serial consoles or KVM switches. However, if the external communication to the rack PDU is just not available, the design of the rack PDU should ensure that basic power distribution, as well as the operation of local management modes such as onboard display, is not affected. It is important that rack PDUs have an automated management path that will maintain basic power distribution. An automated management path also will ensure that when input power to one of the phases of a three-phase intelligent rack PDU is lost, the outlets connected to the unaffected phases continue to be powered.

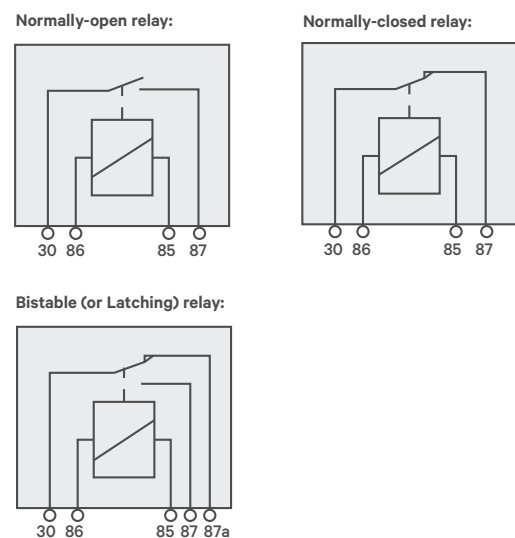


Figure 4: Bistable (or latching) relays resort to providing basic power distribution in the event of a power supply failure.

Maintainability

As computing demands and complexity in the data center continue to rise, unplanned data center outages remain a significant threat to organizations in terms of business disruption, lost revenue and damaged reputation. A 2013 survey of U.S.-based data center professionals by the Ponemon Institute and sponsored by Emerson Network Power showed that an overwhelming majority of respondents had experienced an unplanned data center outage in the past 24 months (91 percent).




Figure 5: Frequency of data center outages over two years, based on results of 2013 survey from the Ponemon Institute.

Regarding the frequency of outages, respondents experienced an average of two complete data center outages during the past two years. Partial outages, or those limited to certain racks, occurred six times in the same timeframe. According to survey responses, complete outages lasted an average of 107 minutes and partial outages lasted an average of 152 minutes. (Figure 5). The second part of the study quantified the cost of an unplanned data center outage at slightly more than \$7,900 per minute.

While the survey provides good data for a more broad discussion on data center downtime and the steps that can be taken to increase availability, it does underline the importance of minimizing the Mean Time To Repair (MTTR) if an intelligent rack PDU goes down.

There are three factors that need to be considered for better maintainability.

- 1. Overcurrent Protection Type:** As mentioned earlier in this white paper, circuit breakers are usually resettable when they trip; while fuses usually require replacement. The procedure for replacing fuses typically requires an electrician and requires the input circuit to be turned off while it is being replaced. Not only does this take time, it requires coordination with facilities teams.
- 2. Modularity:** Modularity along with hot swappability ensures a faster MTTR. Look for a design with a modular communications card which can be replaced while the unit remains operational and providing basic power distribution. Some rack PDUs also provide modularity of input and output power.

- 3. Switching:** Switching capability within a rack PDU ensures that if connected IT gear hangs up, power can be remotely turned on, turned off, or recycled without any physical intervention at the data center. In order to ensure that power to the appropriate equipment is being recycled, care should be taken while associating rack PDU outlets to IT gear. Rack PDUs that provide integration with data center IT equipment access and control solutions simplify such association and minimize the opportunity for errors.

Adaptability

As a result of the rapid pace of change that is occurring in the data center environment, one of the biggest challenges today's data centers face is meeting current requirements while ensuring the ability to adapt to future demands. In the past, this was accomplished by oversizing infrastructure systems and letting the data center grow into its infrastructure over time. Many data centers are abandoning the oversizing approach because it is inefficient in terms of both capital and energy costs.

This also holds true at the rack level. Rack-level flexibility is an important factor in helping data centers adapt to constant change, which often means higher densities and the demand for greater efficiency and control. The need for change can be caused by consolidation, moving from one server or network to another server or network, or the addition of new equipment.

For instance, to handle higher rack densities, there is a need to seamlessly make the change while protecting the existing investment. Intelligent, adaptive rack PDUs with separate power entry modules provide the ability to quickly react to

this change. In addition, hot-swappable modular output power ensures that you protect initial investment and minimize downtime when the server architecture within the rack changes. A modular busway system, which runs across the top of the row or below the raised floor, can also be used to support change by providing the flexibility to add or modify rack layouts and change receptacle requirements without risking power system down time. (Figure 6).




Figure 6: A modular busway system, seen here running across the top of the rows, can be used to support power distribution to the rack.

Conclusion

With increasing data center densities, a single rack can now support the same computing capacity that used to require an entire room. Visibility into conditions in the rack can help prevent many of the most common threats to rack-based equipment, including interruptions to the power source, accidental or malicious tampering, and the presence of water, smoke and excess humidity or temperature.

Today's intelligent rack PDUs can provide the best view of IT power consumption inside the rack. They also can provide a level of monitoring and control capabilities that were unheard of just a few short years ago. However, to ensure the technology provides the benefits of a high availability solution, it's important that considerations are given to the design of the rack PDU, the capabilities it offers, its ability to provide faster Mean Time To Repair, and the level of flexibility it provides.

