Structuring Rack Power

for Edge Computing success

Structuring Rack Power for Edge Computing Success

Why it's essential to enlist the help of a rack power expert to help navigate the huge variety of products on the market.

Data center requirements and technology have been evolving at a rapid pace in recent years. During the last few decades, most businesses moved their server usage from in-house to the cloud and pushed the expertise for server computer hardware outside of their own companies.

However, as technology has progressed, the need for maximal operational efficiencies, performance, and safety has also increased. Today's fast-paced and data-rich digital landscape means companies must ensure there are no breaks in uptime and that data storage and transfer is as efficient as possible. This has led to a boom in edge computing.

When it comes to deploying rack power in edge data centers, the specifications set out by the Open Compute Project are widely misunderstood. That's why it's essential to enlist the help of a rack power expert to help navigate the huge variety of products on the market and the options for deploying them.

This white paper will cover:

- **Why** more companies than ever are making the switch to edge computing
- **The risks** of deploying edge data centers without the help of an expert
- What the **Open Compute Project** is and how companies should approach it
- **The varieties** of OCP rack power configurations and the trade-offs involved
- How to get started in edge computing with **Murata's expert consultancy team**

Understanding Edge Computing and Edge Data Centers

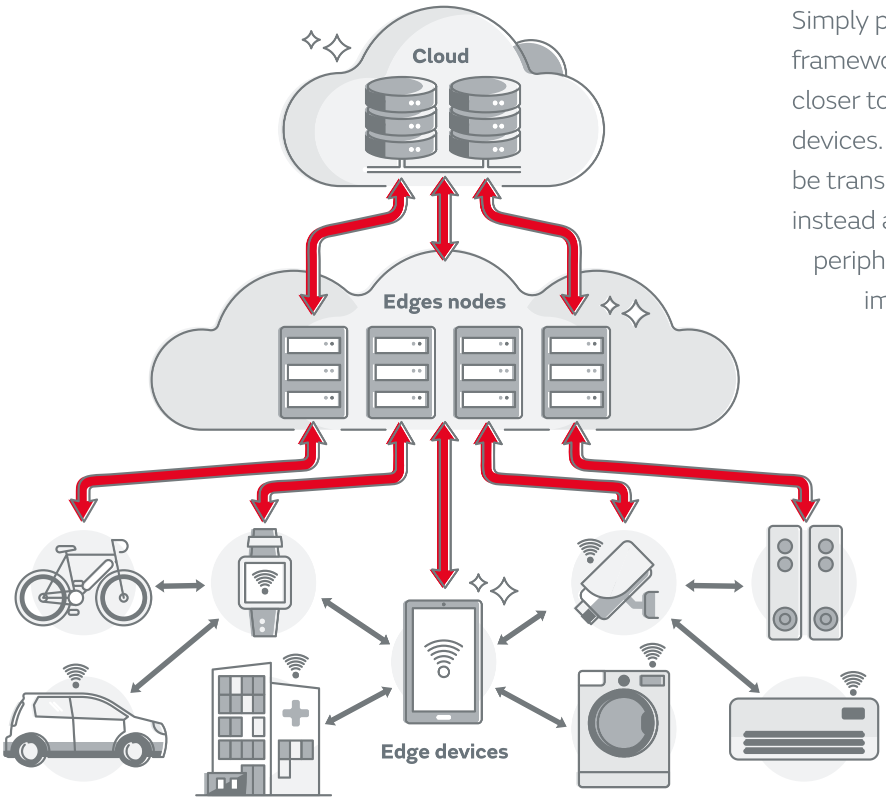

What is Edge Computing?

Simply put, edge computing is a distributed IT framework that moves both computing and storage closer to the sources of data, such as connected devices. It removes the requirement for raw data to be transmitted to a large, centralized data center, instead allowing this data to be processed at the peripheries of the network to reduce latency and improve overall connectivity.

According to research carried out by Markets and Markets, the value of the edge data center market will more than double from $7.2 billion in 2021 to $19.1 billion in 2026. The rapid take-up has been driven by technological advances such as next-generation telecommunications networks (5G), the Internet of Things (IoT), and Augmented and Virtual Reality (AR/VR) – all of which can be better enabled by edge data centers. Indeed, Gartner predicts that by 2025, 75% of enterprise-generated data will be produced outside of a centralized data center.

> ...**the value of the edge data center market will more than double** from $7.2 billion in 2021 **to $19.1 billion in 2026**

Another major incentive for companies to consider bringing their data centers to the edge is the widening data gap. With the exponential and unprecedented increase in volume, dispersing this data through decentralized, low-cost data centers is a logical solution. This makes particular sense in industries such as retail where monitoring inventory is crucial or in data-hungry use cases like factory automation.

The Benefits of Edge Computing

There are good reasons for companies to consider factoring edge computing into their business.

Here are five of the key benefits that edge data centers can deliver:

1. Reduced latency and increased speed

Low-latency, high speed data transfer is one of the major benefits of edge computing. Many modern businesses can't afford to wait for information to travel back and forth to the cloud. Not only can latency be costly, but it can also have drastic implications for the safety and/or functionality of an application. AI in an autonomous application is one example. In this situation, the algorithm relies on huge amounts of data being transferred instantaneously.

2. Increased reliability

Some locations just don't allow for easy transfer and usage of the data to the cloud - and when issues occur, this downtime can be extremely costly. With a centralized data center, periods of downtime can impact every single one of a company's locations and potentially cost tens of millions of dollars in disrupted operations before the issue is fixed. For that reason, certain business sites need to have the capability to store, analyze, and use their data on-site.

3. Increased security

When information travels back and forth to the cloud, there are opportunities for that information to become intercepted. If the information can be utilized closer to its use, it reduces the risk of this type of breach. Additionally, though edge data centers can still be vulnerable to cyberattacks, the amount of data stored in these localized centers is minimal in comparison to the centralized cloud. For this reason, the stakes are much lower if any breaches do occur.

4. Ability to handle valuable information

Not all information has the same value to a business. It makes sense to spend more money securing, managing, and moving data that is more valuable, while handling lower value data in a more cost-effective manner. This is a particularly good strategy to adopt at a time when the cost of moving data back and forth from the cloud - and the bandwidth required to do so - are set to rise significantly.

5. Increased scalability

Centralized cloud data centers can be an incredibly useful asset, but it can be both expensive and labor-intensive to make changes to these structures. By moving computing and storage power closer to the source, infrastructure can be adapted and scaled as required. The modular design of edge data centers allows for increased flexibility and the ability to make fine-grain modifications.

For these reasons, many companies are putting plans in place to create their own Edge Computing Data Centers.

> The modular design of edge data centers allows for increased flexibility and the ability to make fine-grain modifications.

Edge Data Centers:
Getting Started with Edge Computing

Edge data centers, or micro data centers (MDCs) are the smaller versions of traditional data centers that are used to power edge computing.

These data centers contain much of the same architecture that comprises a centralized data center - including power infrastructure, compute, cooling and storage - but their compact size makes it possible to deploy them in a more localized setting.

Some companies may decide to build a network of MDCs to power edge computing across all sites, for example in every single branch of a retail chain. Others may consider rolling out edge computing more gradually or combining it with a centralized cloud as part of a hybrid strategy.

Ultimately, any business looking to get started with edge computing will need to start with the design and deployment of at least one micro data center.

Low-latency, high speed data transfer is one of the major benefits of edge computing.

Edge data centers give the capability to store, analyze, and use your data on-site.

The modular design of edge data centers allows for increased scalability.

Edge data centers provide increased security compared to large data storage facilities.

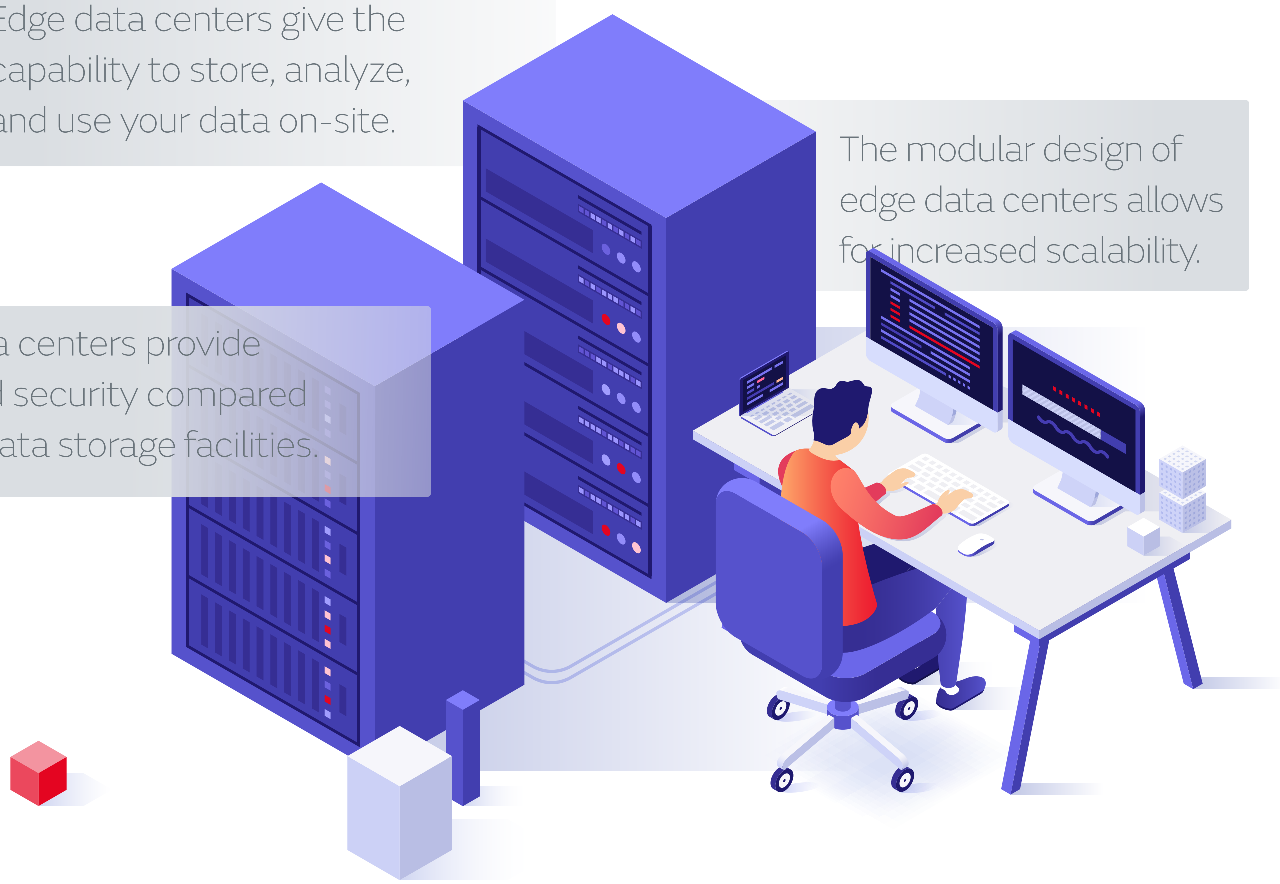

Why Should Companies work with a Rack Power Expert?

Most companies don't have the benefit of an in-house power expert to help them decide how to structure and put in place the power architecture in their edge compute system. However, when companies attempt this intricate task without expert assistance, it can be much more difficult to deploy an edge data center that fully suits their needs.

Here are four of the key reasons businesses are turning to rack power consultants to assist them with their infrastructure.

1. Complexity of edge computing

Deciding what power structure to use in some of today's power-hungry servers can be a daunting task. Companies need to be able to understand the trade-offs and variables involved in designing the optimum rack power structure and ask the right questions to avoid costly mistakes in implementation.

2. Undertaking an informed cost/benefit analysis

Decentralized edge data centers may be simpler and more cost-effective than centrally located data centers, but when multiple edge data centers are required, the complexity of this infrastructure increases exponentially.

As PriceWaterhouseCoopers asserted in a recent study: "If not carefully built and managed, the complexities of managing and securing such a network could dilute or even overwhelm the potential cost savings and efficiencies."

To avoid this happening, an expert can walk you through a realistic cost/benefit analysis of various edge data center setups and assist you in adopting a cost-effective and future-proof approach.

..."If not carefully built and managed, the complexities of managing and securing such a network could dilute or even overwhelm the potential cost savings and efficiencies."

3. There is no "one-size fits all" approach

While open-source specifications are available on the market, these should be considered a guideline rather than a firm template. Companies need to understand how different rack power setups can apply to their specific needs and priorities, whether that involves connecting to thousands of IoT devices, ensuring continuous uptime, or processing an ever-increasing wealth of data with minimal latency.

4. The ongoing cost of poor rack-power design

There are several variables at play when it comes to deciding on the right rack power setup - and making the wrong choice could cause issues in the future. One key example of this is deciding on the required voltage, which can range anywhere between 12V and 54V. While picking a lower voltage can initially save money, attempting to increase power at a later date is problematic and pricey. Conversely, overestimating the power required can lead to unnecessary cost.

muRata
INNOVATOR IN ELECTRONICS

Variations & specifications for Rack Power Structures

Understanding the Open Compute Project

Though rack power in edge data centers can be complex, existing specifications and frameworks from the Open Compute Project Foundation (OCP) have eliminated the need to start from scratch.

> Established in 2011, the **OCP Foundation** was founded to promote open-source and collaborative hardware just like software had done in the past.

Established in 2011, the OCP Foundation was founded to promote open-source and collaborative hardware just like software had done in the past.

The hope was for this to rapidly increase the pace of innovation and development in the data center hardware arena in a way that was accessible to all. Over the years, this OCP architecture has taken hold and become a standard - albeit with several customizable options.

It's important to note that OCP's open-source specifications for hardware in rack power are guidelines that can be modified and adapted, for example, to ensure a more reliable power supply or lower operational costs.

Designing a well-optimized rack power setup involves numerous components that need to be considered, including:

- **Output voltage selection**
- **Output power calculation**
- **Input power source options**
- **Battery backup options**
- **Thermal management**
- **High-performance servers**
- **Automatic Transfer Switch (ATS) considerations**

By detailing some of the various iterations of OCP's rack power specifications, we hope to give you a sense of the sheer number of variations and options on the market.

Variations & specifications continued...

Open Rack Version 1 (ORV 1)

The Open Rack Version 1, ORV1 or V1, OCP utilized a 12V, three-busbar approach in a 21" rack with 1 to 3 power zones.

The input power was specified to be three phase or single phase with up to 277VAC nominal input, single phase to each power supply.

Each power zone had a power shelf and there were three busbar pairs coming out of the back of the power shelf that slotted into the 21" rack busbars. Each power shelf could be loaded with up to (7) 700W OCP power supplies with maximum dimensions of 68mm X 125mm X 300mm. The power was specified to be up to 4200W N+1 per shelf.

The ORV1 utilized a separate rack for the 12V battery backup that supplied backup power when the AC source went away until the generators could kick in.

ORV1 has since been replaced by ORV2.

Open Rack Version 2 (ORV 2)

The Open Rack Version 2, ORV2 or V2 utilizes a 12V single busbar approach with two separate power zones.

The ORV2 rack specification also allows for a single power zone and or three busbars if desired. There are a lot of other customizable options available to the user.

Facebook released an ORV2, 12V power shelf single bus-bar design that included (3) 3.6kW power supplies with three battery backup lithium-ion BBU modules. One BBU worked in conjunction with each power supply. There is a bidirectional DC to DC in each power supply that charges the batteries and reverses to apply 12V on the busbar to support the system for 90 seconds with 12V when the AC sources are down.

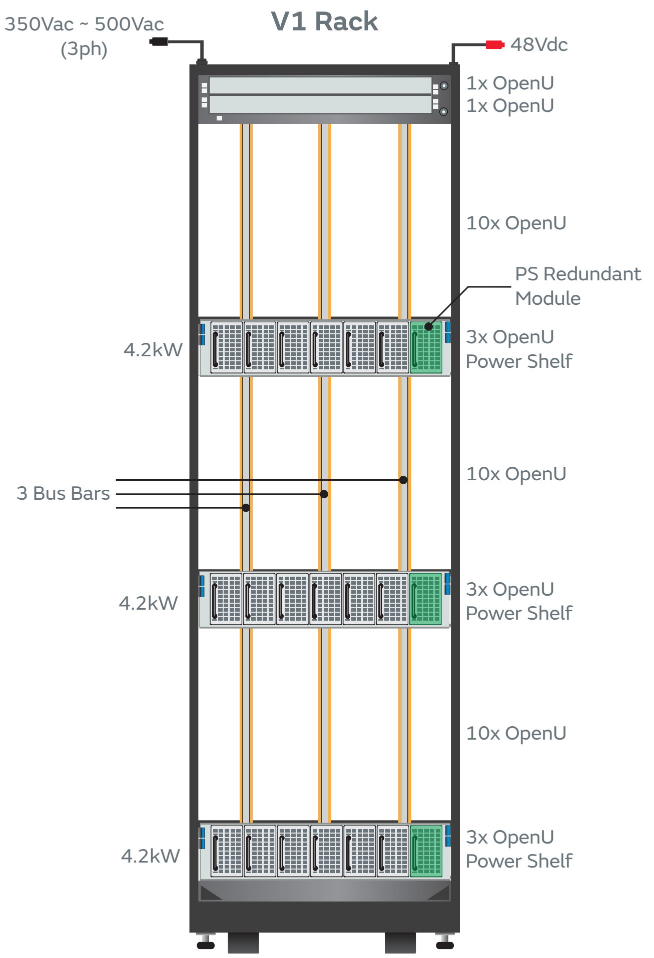

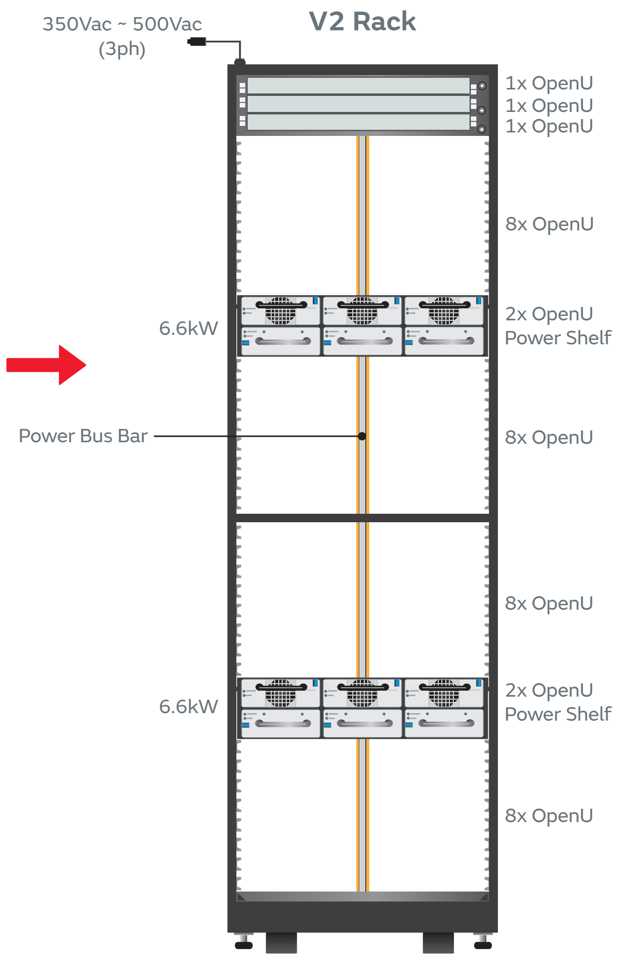

muRata INNOVATOR IN ELECTRONICS

1OU 21” 12V OCP power shelf

Several manufacturers, including Murata, have developed an 1OU 21” 12V OCP power shelf that can be assembled into the ORV2 rack with either one or three busbars.

The 12V busbar is bolted to the rack busbar in the back of the rack. This central rack-power architecture was developed to eliminate the need to change the power every time the server requirements changed. It also used an N+1 power supply scheme instead of the previous N+N power supply requirement in each server, which cut costs by reducing the overall redundancy.

The Murata 21” shelf

In the Murata 21” shelf, up to six 3kW 12V power supplies can be installed in a 15kW N+1 configuration where all the power supplies share the load equally.

Murata also has a 2U BBU shelf that takes power directly from the busbar and charges the batteries during normal operation and then supplies 12V to the system for up to two minutes of backup when the AC is lost.

The Murata 2 RU shelf

Murata has developed a 2 RU shelf that can fit in a 19” EIA rack or 21” OCP style rack.

This power shelf enables the use of shelf-level Automatic Transfer Switches (ATS) that allow for two separate input power sources. During normal operations, AC source one is used. When AC source one is lost, the ATS switches to using AC source two without the output being affected. This enables the use of otherwise redundant input power.

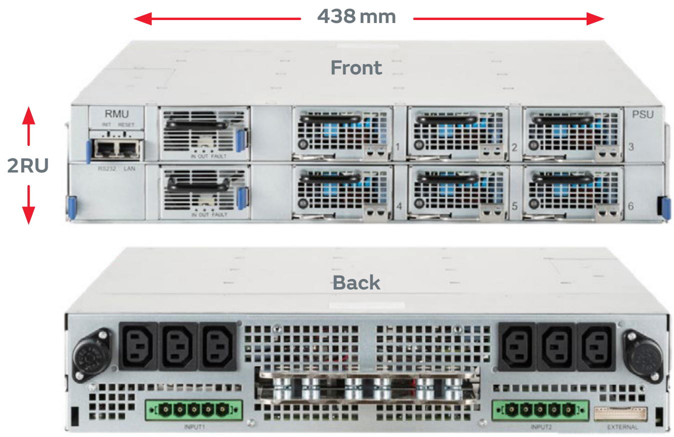

In 2016 Google contributed a 48V ORV2 specification to the Open Compute Project that could be used in conjunction with the ORV2 rack. One of the reasons Google used 48V was because the voltage drop on the busbar is reduced dramatically, increasing the overall efficiency of the system.

In the context of Google’s ORV2 specification, the mention of 48V can be misleading. In fact, this is the telecom 48V, which means a typical 54V system. That comes from the 48V stack of lead acid battery cells that equates to 54V when fully charged.

The Murata six-across 54V 1OU OCP power shelf

Murata has a six-across 54V 1OU OCP power shelf with a flexible output busbar configuration. This shelf provides 18kW in an N+1 configuration.

This 54V shelf can fit in a 19” rack or with an adapter kit fit into a 21” OCP rack.

Open Rack Version 3 (ORV 3)

When the OCP team approached the ORV3 specification, they originally felt the 48V system was what they wanted to use to achieve more power density in the rack and increase efficiency.

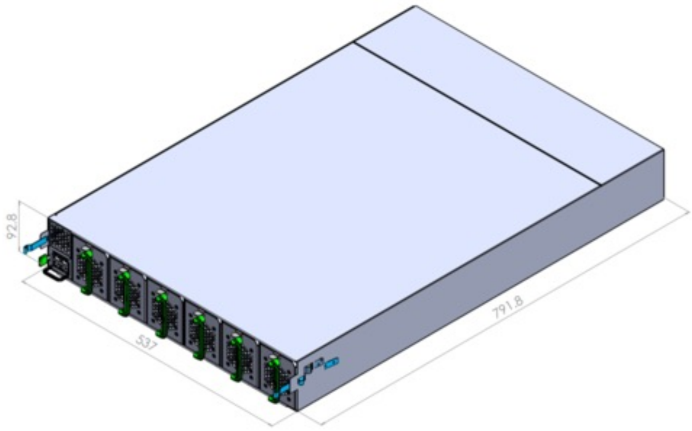

During their initial testing, they opted for a 50V regulated output. By using 50V, they could then use a 4 to 1 DC to DC converter instead of a 5 to 1 DC to DC to provide 12V in the server. This slightly increased the efficiency in the transition from 50V to 12V vs 54V to 12V in the server and therefore led to a decrease in operational costs.

They eventually raised the output voltage by 1V to 51V out to allow for more time to turn on the ORV3 BBU when AC is lost. The ORV3 BBU transitions to backup mode when the busbar voltage reaches 47.5V. This still allowed for the 4 to 1 DC to DC converter in the server.

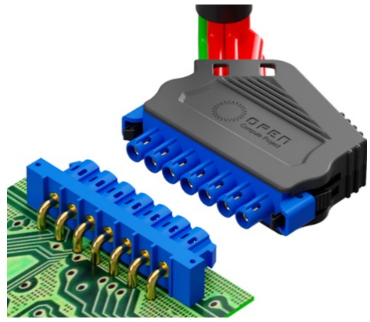

The ORV3 BBU has its own DC to DC converter that charges its lithium-ion batteries from its connection to the busbar. When the BBU sees 47.5V on the busbar the internal DC to DC supplies regulated 48V back to the busbar, to back-up the rack for at least four minutes at full power.

The ORV3 system uses a specific input connector that allows for different input voltage configurations by changing the AC whip.

The ORV3 power shelf and BBU output with 51V and 48V respectively is connected to a special ORV3 busbar via a press fit ORV3 power connector. The press fit power connector on the output of the power shelf is like the previous IT gear press fit connector used for IT gear in the ORV2 system. Because the current is just a quarter of that provided by a 12V system, a press fit connector is viable and makes the assembly of the shelf to the busbar less cumbersome.

Another reason for the BBU disaggregated from the power supply in a standalone module is that it allows the BBU to monitor the DC bus directly. This makes for a smoother transition than the previous ORV2 system that monitored the AC input to the power supply and transitioned when it was determined the AC went away.

Murata's latest rack power products

The Murata subsystems team is currently developing a version of the ORV3 power shelf with six-across 1OU power supplies. This is set to be released during Q4, 2022. The shelf will be identical to the OCP version. The company will use its 3.6kW, 54V power supply modified to 51V at 3.6kW, optimized for high efficiency. The OCP ORV3 specification calls for a 3kW power supply.

Murata will also have a 48V version of the 3.6kW power supply. The 54V, 51V, or 48V power supplies will be able to fit in any of the Murata shelves designed to house the 68mm power supplies.

Murata is presently developing the ORV3 BBU to meet all the ORV3 BBU specifications. The ORV3 BBU is scheduled to be available in late 2023.

How to avoid costly implementation mistakes

Enlist the help of a world-class team

Murata Manufacturing is known for being the leading capacitor company in the world, but it also offers deep expertise in both the manufacture and implementation of rack power products.

The Murata Power Solutions division is responsible for designing and manufacturing power supplies, while the Murata Subsystems team designs and manufactures OCP, hyperscale, and rack power.

The OCP rack power products are developed in Hong Kong and Japan by two world-leading design teams with centuries of combined experience.

The Hong Kong team comprises of 21 engineers with over 420 years of power experience as well as an additional five members of support staff.

Meanwhile, the Japan design team is composed of 31 engineers with over 500 years of experience and supported by an additional 27 engineers in Japan.

Hong Kong team comprises of **21 engineers** with over 420 years of power experience

The **Japan design team** is composed of **31 engineers** with over 500 years of experience

Murata's OCP and rack power consultants

Murata has several leading consultants on board who are available to assist with rack power design and implementation issues.

Kevin Gero is the Senior Product Marketing Manager based in North America. Kevin has been in the power industry for over 40 years at both Murata and other companies working with Facebook and has been involved in the Open Compute Project since its inception.

Taz Suzuki is an Applications Engineer at Murata and has been a key member of the engineering team in Japan working on the controller used in the shelf for several years. Taz is now in San Jose supporting Murata's OCP customers.

With the entire subsystems design team behind them, Taz and Kevin are well-placed to answer almost any question regarding power.

Taking the First Step in Edge Computing

From higher speeds to lower costs, the benefits of moving compute power and storage closer to the edge are innumerable. To unlock these benefits, opting for the right configuration of OCP power supplies is essential - but with numerous variables and trade-offs, navigating through the differences can be a difficult task.

Leveraging our OCP and power optimization expertise, Murata's consultants are here to help.

To find out more about Murata's rack power consultancy services and products, contact us today.

Contact: kevin.gero@murata.com

Sources

1. https://www.marketsandmarkets.com/Market-Reports/edge-data-center-market-142018469.html
2. https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders
3. https://www.pwc.com/us/en/industries/capital-projects-infrastructure/library/assets/pwc-edge-data-centers.pdf
4. https://datacenterfrontier.com/edge-computing-use-cases-from-ai-and-telehealth-to-factories-and-5g/
5. https://www.techtarget.com/iotagenda/tip/Top-5-benefits-of-edge-computing-for-businesses
6. https://www.pwc.com/us/en/industries/capital-projects-infrastructure/library/assets/pwc-edge-data-centers.pdf

muRata
INNOVATOR IN ELECTRONICS