molex

Setting the Bar for Data Center Architecture

Molex and the Open Compute Project

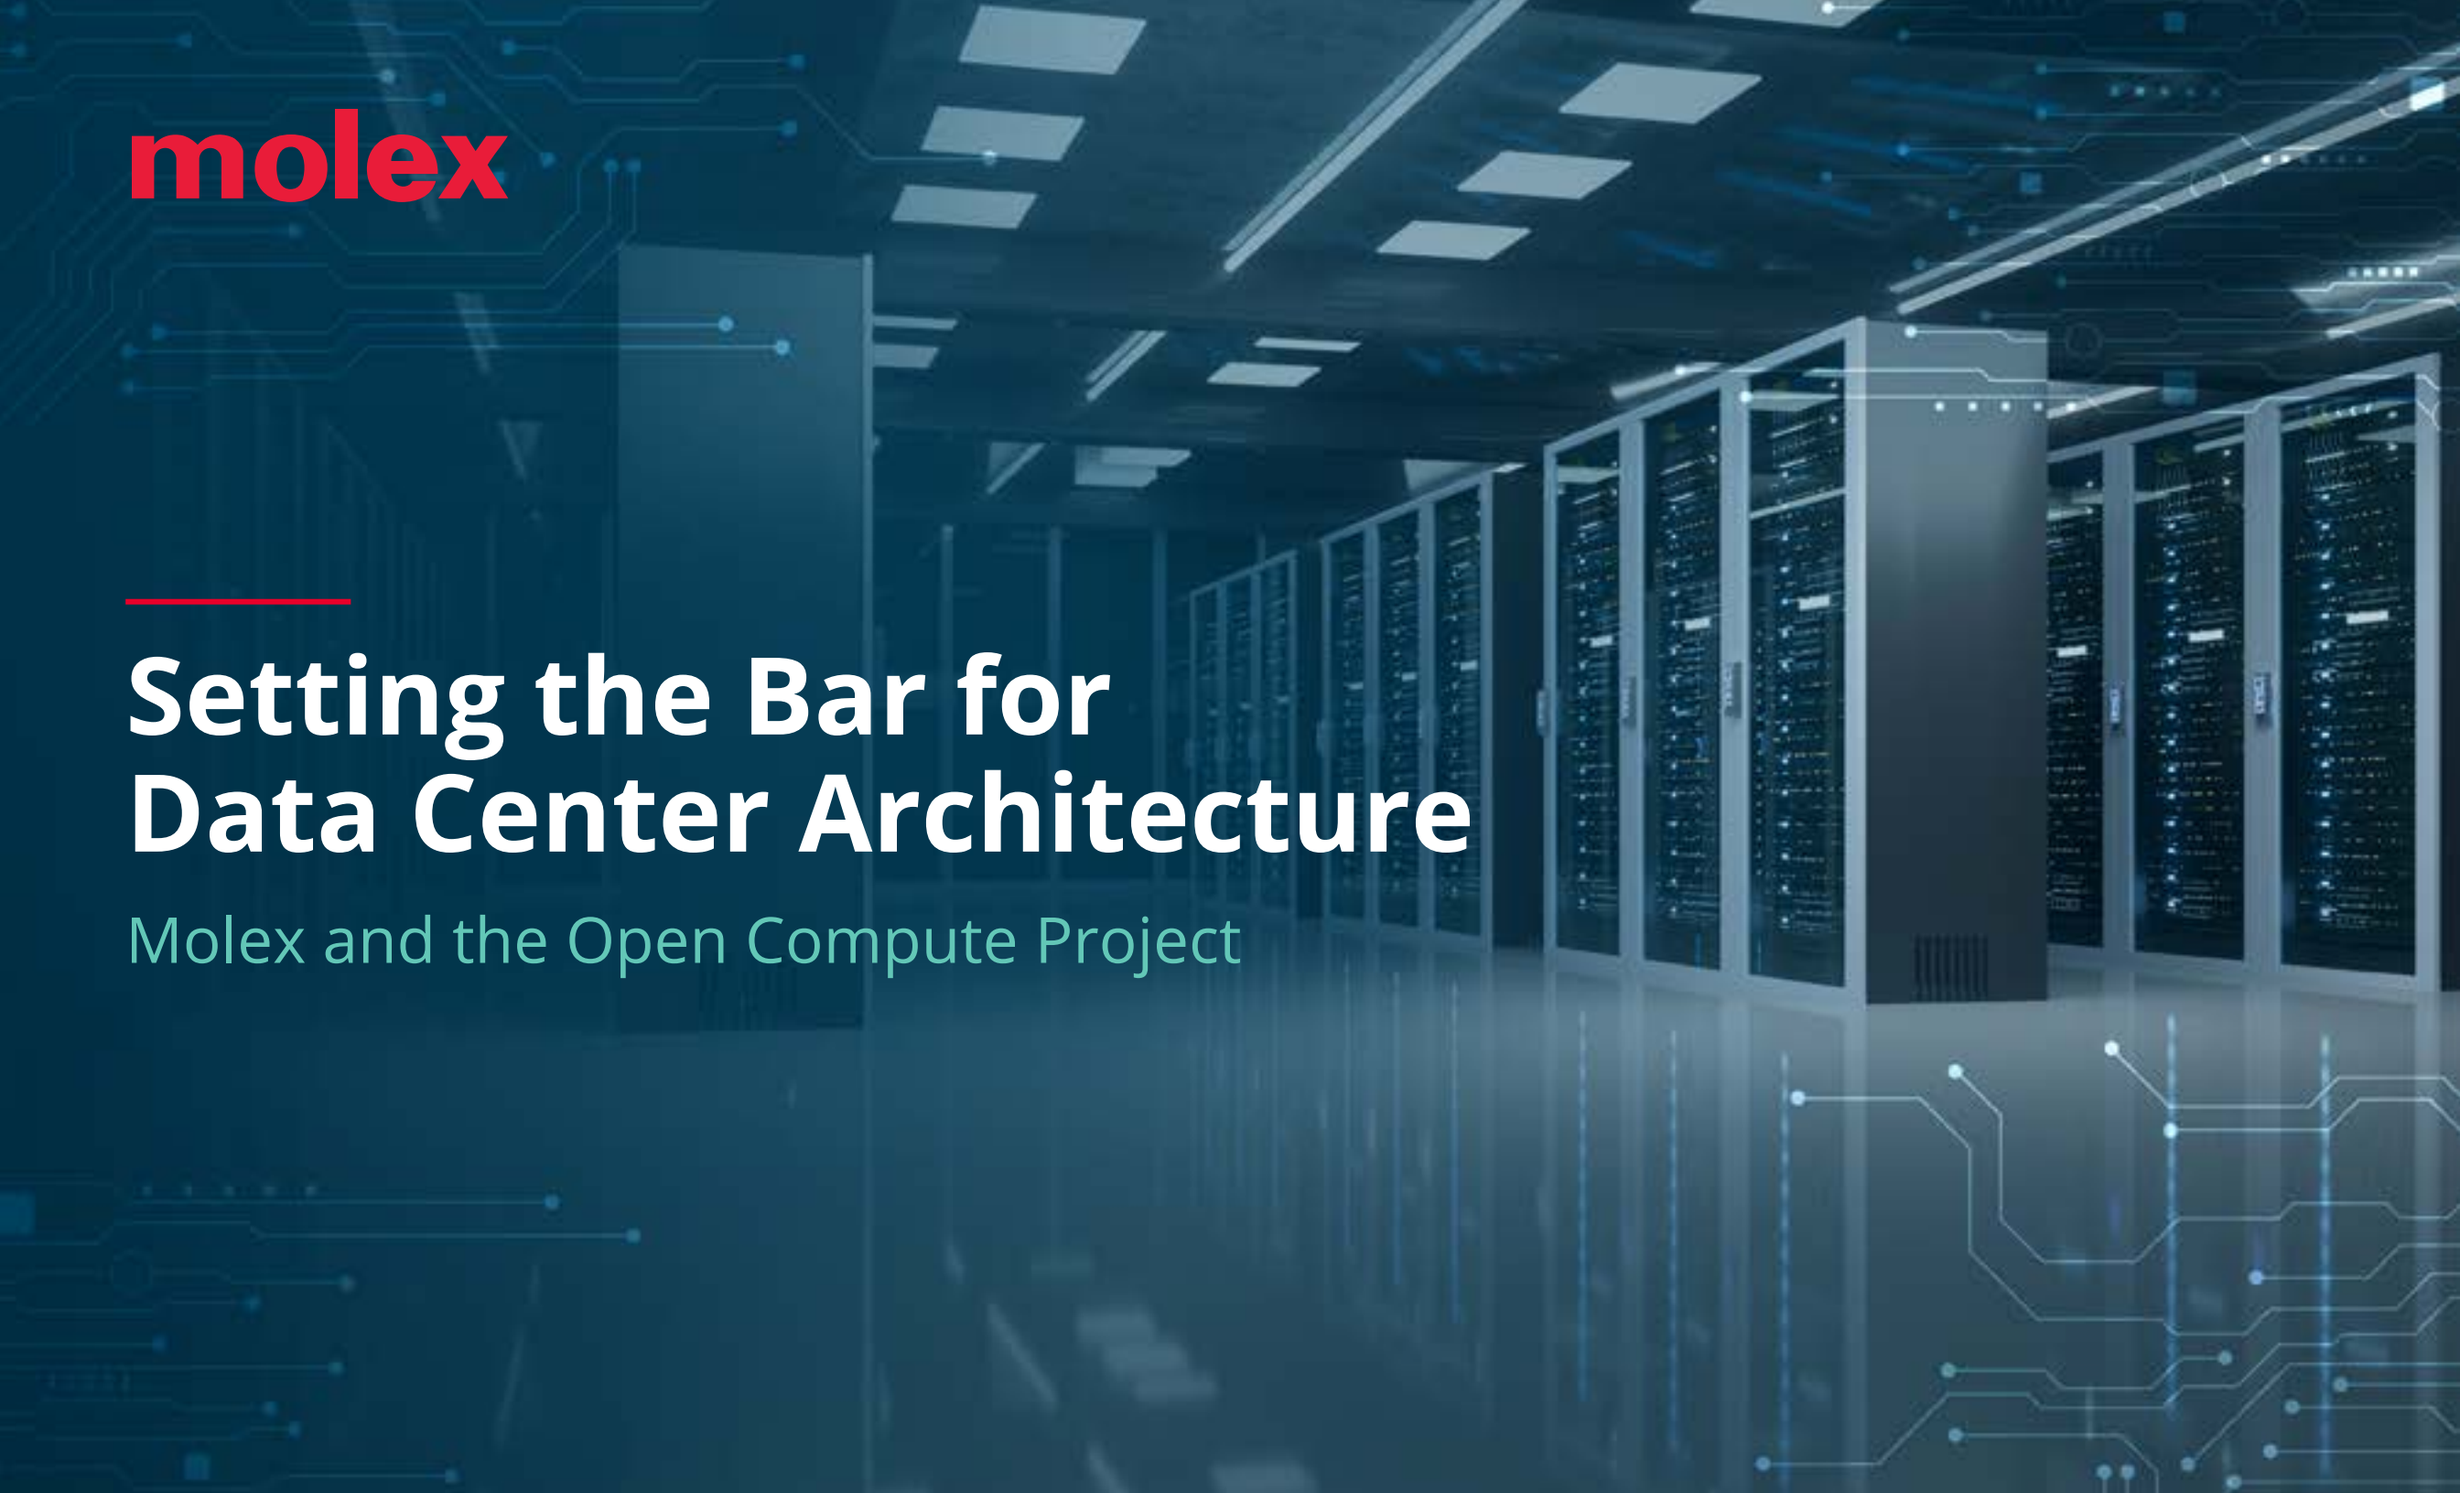

OCP: BUILDING THE BLUEPRINT FOR NEXT-GEN DATA CENTERS

Optimal efficiency. Seamless standardization. Maximum performance. These are the resounding themes of the **Open Compute Project** (OCP). As the industry transitions toward disaggregated architectures, OCP is leading the charge, defining the future of data centers through open standards and collaborative innovation. OCP's impact is broad, spanning the entire data center ecosystem — from storage and networking to emerging technologies like CXL and SONiC. This comprehensive strategy ensures that data centers remain adaptable to evolving technologies and business needs.

OCP's emphasis on efficiency, standardization and performance aligns seamlessly with the industry's shift toward disaggregated architectures. This approach accelerates time-to-market, reduces operational costs and enhances scalability. Key OCP initiatives like DC-MHS, OAI and ORV3 Rack and Power are on the front lines of this transformation, setting standards for mechanical specifications, accelerator integration and power infrastructure.

Molex, a dedicated OCP member, is at the forefront of developing high-density interconnect solutions that empower hyperscale success. Our expertise in connector design and manufacturing, combined with a deep understanding of industry demands, allows us to deliver products that exceed industry standards and provide tangible business value.

By actively participating in OCP, Molex contributes to the creation of reference designs that serve as blueprints for efficient and scalable data center infrastructure. Our high-density connectors are engineered to meet the rigorous requirements of these designs, ensuring superior performance, reliability and interoperability. Beyond product development, Molex engages in industry forums and collaborates with fellow OCP members to shape the future of data center connectivity. Our mission is to drive innovation, reduce costs and enhance overall data center efficiency.

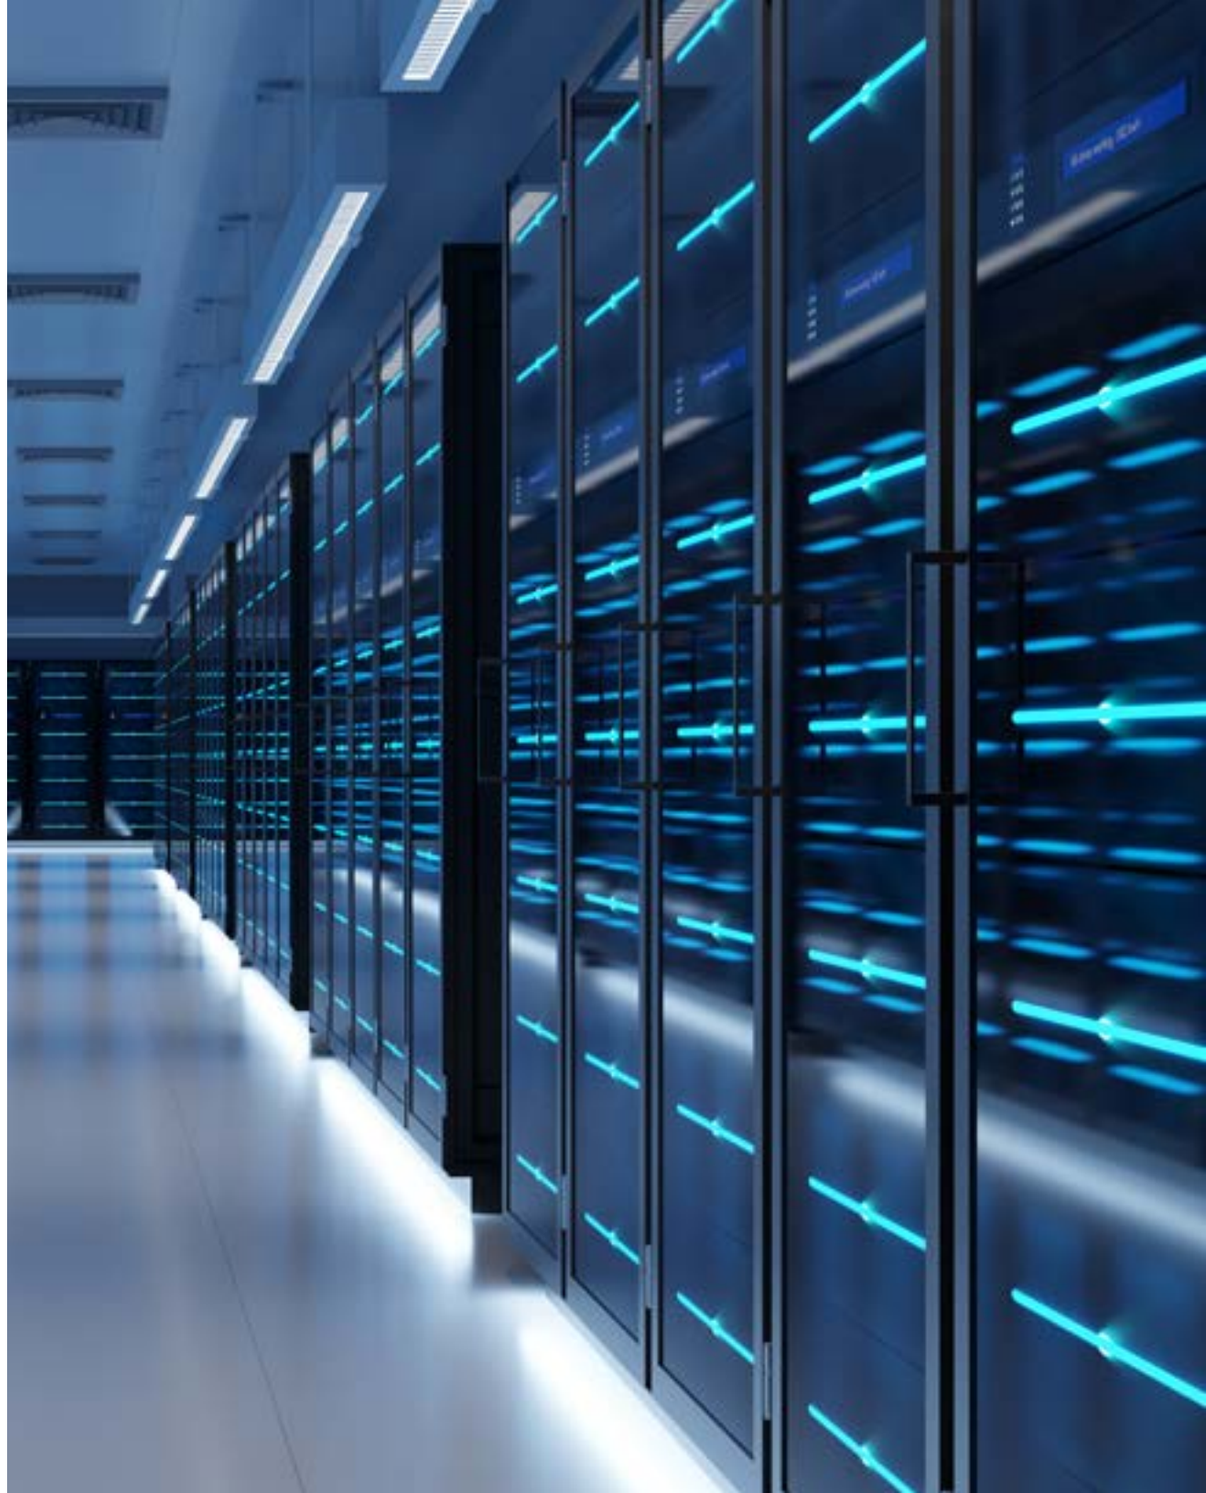

DC-MHS: STANDARDIZING EFFICIENCY AND SCALABILITY

The **Data Center Modular Hardware System** (DC-MHS) specification provides a unified framework for optimizing data center performance and adaptability. DC-MHS establishes a foundation for modular, interoperable and adaptable data centers by creating protocols for physical infrastructure components. This specification covers critical elements such as racks, power distribution units and cooling systems.

Addressing key challenges in power delivery, thermal management and cable management, DC-MHS enables data center operators to build cost-effective infrastructures capable of supporting evolving workloads, ensuring readiness for future demands.

One of the critical advancements furthered by DC-MHS is the support for PCIe Gen6, a crucial component as the industry moves towards higher data rates and increased bandwidth. By standardizing the physical infrastructure and ensuring seamless integration, DC-MHS positions data centers to fully leverage PCIe Gen6 capabilities, which are essential for next-generation applications. This focus not only prepares data centers for the current demands but also ensures scalability for future innovations, solidifying DC-MHS as a cornerstone in the evolution of data center architecture.

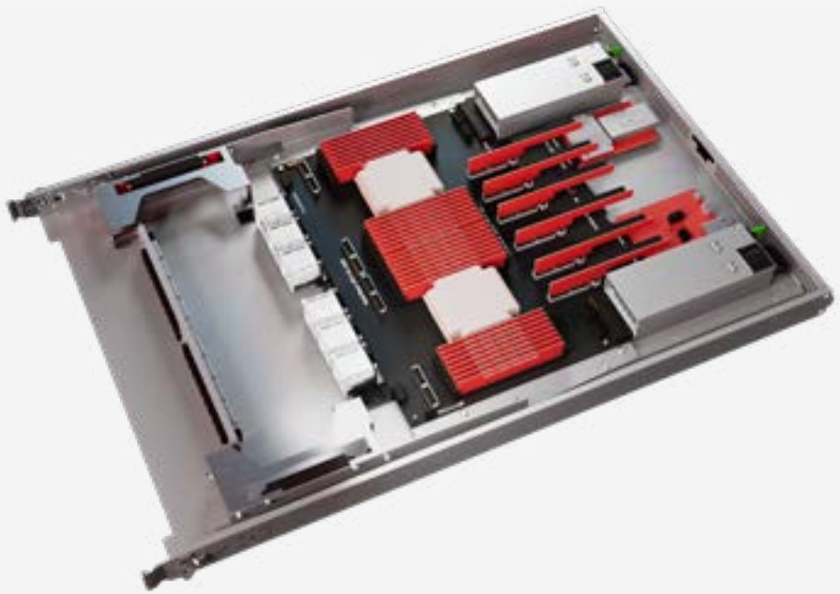

KEY CHALLENGES IN DC-MHS IMPLEMENTATION

High-Density Power and Cooling Requirements

The push for greater computational power is driving an increase in the concentration of servers within data centers, leading to higher power demands. This growth necessitates robust power delivery systems and, therefore, advanced cooling strategies. This puts a bright spotlight on the need to address head-on the significant challenge of balancing energy efficiency with the need for high performance.

Efficient Cable Management and Routing

Managing the complex network of cables in high-density racks is crucial. Optimizing cable routing, reducing congestion and ensuring proper airflow are not simple undertakings but are essential to preventing equipment failures and maintaining optimal system performance.

Integration of Various Components

With the vast range of components, including servers, storage and networking equipment that makes up a data center, seamless integration of these elements, while maximizing space and minimizing operational complexity, is critical for efficient operations.

Aligning with Common Connectors and Interfaces

Inconsistencies in connectors and interfaces can impede scalability and interoperability. Utilizing a uniform set of connectors for power, data and signal transmission simplifies system integration and reduces overall costs.

DC-MHS WORKSTREAMS & MOLEX SOLUTIONS

The DC-MHS subproject includes multiple workstreams, each targeting key data center infrastructure components. From peripheral device connectivity to power distribution and shared infrastructure, these workstreams focus on enhancing efficiency and scalability.

M-XIO/PESTI: Peripheral Device Connectivity

The M-XIO/PESTI (Modular Extensible Input/Output, Peripheral Sideband Tunneling Interface) workstream standardizes peripheral device connectivity in data centers, addressing challenges like high-speed data transfer and signal integrity. The direct-to-contact termination strategy eliminates the paddle card from the cable assembly, enabling more repeatable and reliable automation and creating a direct connection from anywhere in the system to near the ASIC. Molex's **NearStack PCIe Connector System** provides a high-performance solution, delivering the bandwidth and reliability required for demanding peripheral applications. With a small pitch and reduced mated height, the NearStack PCIe Connector takes up minimal space within the system, alleviating space constraints.

M-PIC: Host Processor Module Connectivity

M-PIC (Modular Platform Infrastructure Connectivity) focuses on creating a uniform model for infrastructure connectivity for Host Processor Modules (HPMs), tackling issues related to power delivery, signal integrity and space constraints. Molex's **KickStart Connector System** offers a unique solution that combines signal and power to deliver a robust, scalable solution with a compact, reliable design. Molex's **Micro-Fit+** and **Pico-Clasp** connectors provide versatile options for various power and signal interconnections within the HPM and chassis. As an all-in-one system, KickStart is the first OCP-compliant solution that combines low-speed and high-speed signals — as well as power circuits — into a single cable assembly.

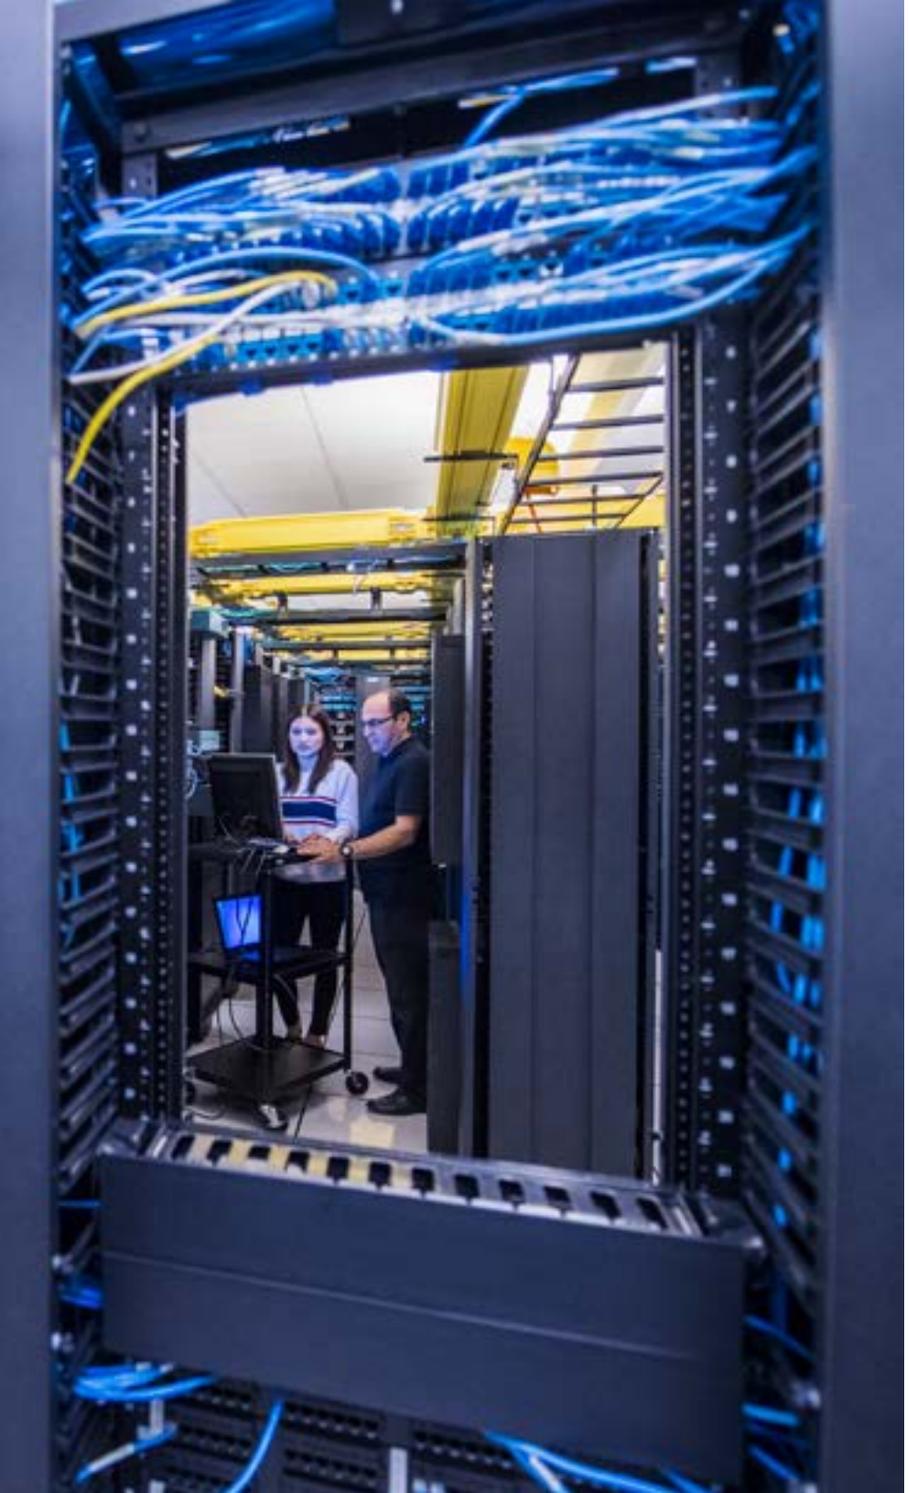

M-CRPS: Power System Reliability

M-CRPS (Modular Common Redundant Power Supply) aims to enhance power system reliability and efficiency by formally systemizing redundant power supplies. This workstream addresses power density, redundancy and thermal management challenges. Molex is actively developing future-proof connector families to meet the stringent requirements of M-CRPS.

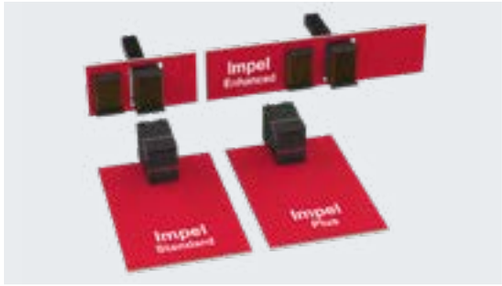

M-SIF: Shared Infrastructure Optimization

M-SIF (Modular Shared Infrastructure) seeks to agree on shared infrastructure components within data center enclosures, focusing on space maximization, power delivery and signal integrity. Molex's Impel Enhanced connector, specified in the OCP M-SIF Base Specification, provides a high-speed, high-density solution to these challenges.

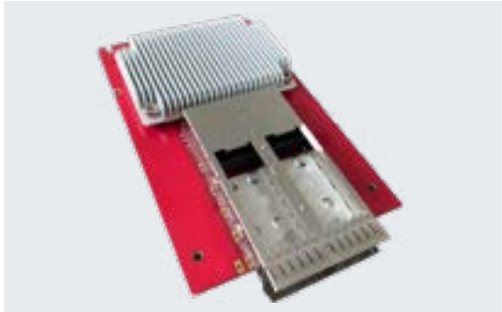

NIC and DC-SCM: High-Speed Networking and Storage

While not a standalone workstream within DC-MHS, network interface cards (NICs) and data center storage and compute modules (DC-SCMs) are critical components significantly impacted by DC-MHS guidelines. These components benefit from a uniform approach to optimizing integration, power delivery and thermal management within the data center environment. Molex's **QSFP-DD** and **Silver 4C+ Edge Card Connectors** deliver high-performance hardware for these critical applications.

	Data Rate		Standard	Mated Height
	PCIe Gen 5 32 Gbps	PCIe Gen 6 64 Gbps		
NearStack	X	X	SFF-TA-1026	9.8 mm
KickStart	X	X	SFF-TA-1036	11 mm

Additional Molex DC-MHS Solutions:

Micro-Fit+ Connectors
These robust and versatile connectors provide a comprehensive range of wire-to-board and wire-to-wire components for power and signal applications in data centers and industrial environments.

M-CRPS Connectors
Designed specifically for redundant power supply systems, M-CRPS Connectors ensure reliable and efficient power distribution in data center infrastructure.

QSFP-DD Connector System
A high-speed, low-profile product portfolio that supports multiple data channels, the QSFP-DD Connector System is ideal for data center applications that require high bandwidth and low latency.

Multi-Trak Connectors
This versatile connector system offers a modular and scalable option for various I/O applications, including power, signal and data transmission.

OPEN ACCELERATOR INFRASTRUCTURE: A CATALYST FOR DATA CENTER TRANSFORMATION

The **Open Accelerator Infrastructure** (OAI) framework aims to formalize accelerator platforms, simplifying the integration of diverse technologies into data centers. By establishing a common infrastructure specification, including universal baseboards and predetermined form factors, OAI allows data center operators to incorporate various accelerator technologies seamlessly. This standardization encourages innovation, reduces development costs and enhances operational efficiency.

Beyond hardware, OAI emphasizes software compatibility, ensuring that accelerators can function seamlessly with existing operating systems. The comprehensive OAI approach to accelerator integration enables the deployment of heterogeneous accelerator applications, unlocking the full potential of AI, machine learning and high-performance computing workloads.

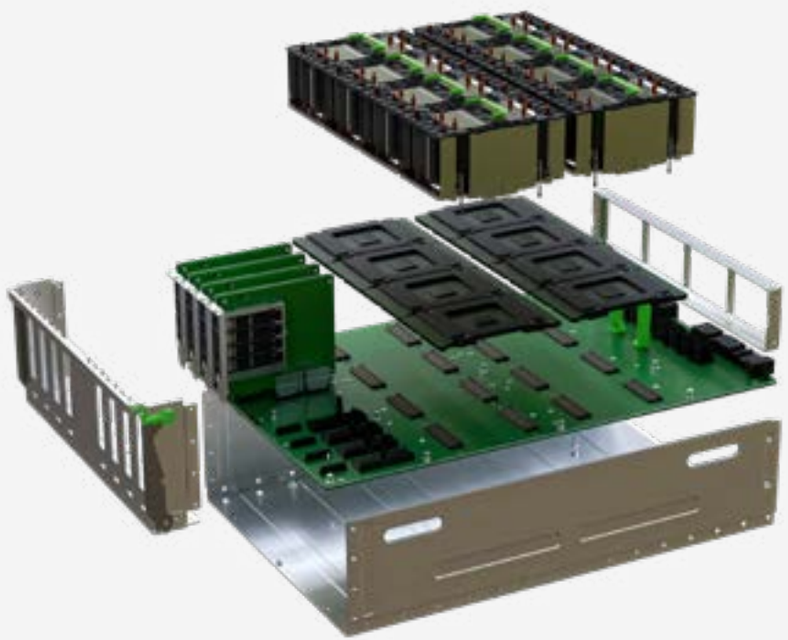

OBSTACLES TO OAI DEPLOYMENT

High-Speed Data Transfer and Low Latency

Accelerators demand exceptionally high bandwidth and minimal latency to elevate their computational potential. Efficient data movement between the accelerator, host CPU and memory is critical for overcoming performance bottlenecks and achieving optimal results. The OAI framework must address these challenges by providing robust interconnect components and optimizing data transfer protocols.

Power Delivery and Thermal Management for Accelerators

Power delivery and thermal management are critical considerations for OAI-based systems. Accelerators often display high power densities, underscoring the need for efficient power distribution networks and advanced thermal strategies. Balancing the need for sufficient power with thermal constraints is essential to preventing system instability and performance degradation. Additionally, accelerators' dynamic power consumption presents challenges for power delivery systems, requiring flexible and responsive power management strategies.

Flexible and Scalable Interconnect Solutions

The dynamic nature of data center environments relies on highly flexible and scalable interconnect components. OAI must address the challenges of connecting diverse accelerator types, varying compute densities and evolving network topologies. Interconnect products must adapt to changing workloads and support high-speed data transfers and scaling to accommodate future growth. The connector infrastructure should also minimize latency and power consumption while ensuring reliability and fault tolerance.

Modular and Upgradable Design

Achieving a truly modular and upgradable OAI architecture calls for careful consideration of several factors. Components must be designed with consistent interfaces to facilitate easy replacement and upgrade. Additionally, the system must support hot-swappable modules to minimize downtime. Balancing flexibility with system integrity is crucial to ensure that upgrades do not compromise overall performance or reliability. Design architects must maintain backward compatibility to protect existing investments and facilitate gradual transitions to newer technologies.

OAI CHALLENGES ADDRESSED BY MOLEX

Mirror Mezz Pro Connectors are at the forefront of OAI interconnect modules. Selected as the OCP reference for the Open Accelerator Module (OAM) v1.5, this connector series delivers exceptional performance, reliability and scalability. Designed for intensive high-speed applications, Mirror Mezz Pro offers superior signal integrity, enabling data rates up to 112 Gbps while consolidating board space through its compact, hermaphroditic design.

Mirror Mezz Enhanced Connectors build upon the success of the Pro series by introducing further advancements for next-generation OAI systems. These connectors offer enhanced performance metrics and expanded capabilities — supporting data rates up to 224 Gbps — making Molex a leader in providing cutting-edge products for the evolving data center landscape.

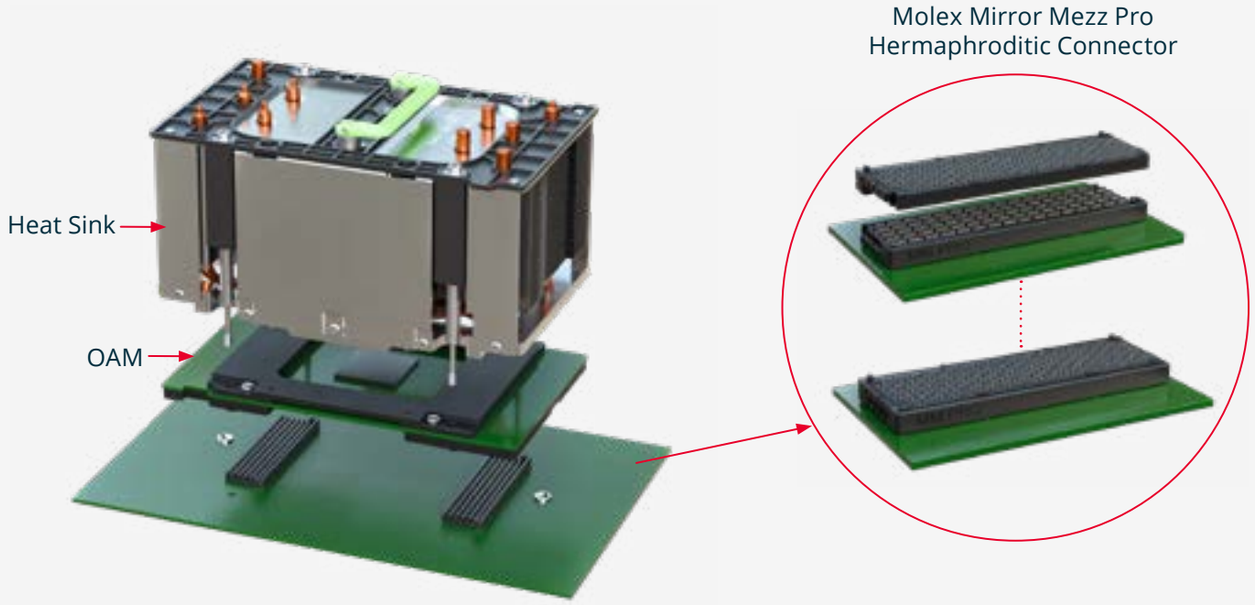

Types of Mirror Mezz	Overall MATED Dimension	Footprint	Full Differential Pairs	Single-Ended Pins	Total Pins
15x11 Standard MM 15x11 MM Pro 	SAME 22.00mm x 68.00mm	Standard MM and MM Pro – same footprint	161	44 SE	688
15x11 Mme (Mirror Mezz Enhanced) 		Different footprint compared to Standard MM and MM Pro	166	24 SE	688

ORV3: A FOUNDATION FOR MODERN DATA CENTERS

ORV3 (Open Rack v3) Rack and Power is a standardized platform designed to improve data center efficiency and scalability. By establishing common specifications for racks, power distribution units (PDUs) and cooling systems, ORV3 empowers data center operators to build flexible and adaptable infrastructure. This hardware uniformity aids interoperability between various components, reducing complexity and accelerating deployment.

ORV3 focuses on strengthening power efficiency, improving airflow and simplifying cable management within the rack environment. Through pre-specified rack dimensions, power delivery and cooling configurations, ORV3 enables data centers to achieve greater server concentration, lower operating costs and improved overall performance.

OVERCOMING ORV3 HURDLES

Increased Server Density and Power Consumption

As data centers push for more computational power, server crowding continues to rise. This expansion warrants greater power consumption, exerting considerable pressure on power delivery systems and thermal management structures. Balancing high performance with energy efficiency becomes a crucial challenge that has traditionally been tough to address.

Efficient Airflow and Thermal Management

Managing heat dissipation in denser racks is a complex task. Effective airflow patterns are key to preventing equipment from overheating and bolstering cooling system performance. Achieving a balance between cooling efficiency and energy consumption is vital for maintaining optimal operating conditions.

Flexible Cabling and Connectivity Components

The evolving landscape of data center infrastructure requires adaptable cabling and connectivity solutions. These must accommodate various server form factors, power demands and network topologies, while minimizing cable congestion and improving space utilization. This presents significant challenges for design and deployment.

Integration of Different Server Form Factors

Data centers often house diverse server form factors, each with unique power, cooling and cabling needs. Integrating these varied systems into a cohesive rack environment while maintaining optimal performance and efficiency is a complex undertaking. Formalizing rack layouts and power distribution while accommodating different form factors helps drive efficient operations.

MOLEX SOLUTIONS FOR ORV3 SUCCESS

Molex **PowerPlane Busbar Connectors** are just one option that addresses the challenges posed by increasing server density and power consumption in ORV3 environments. PowerPlane significantly reduces power losses and improves overall system efficiency by providing a high-current, smaller form factor power distribution solution. The connectors' compact design maximizes available rack space, while its modular architecture sets the stage for easy scalability to accommodate growing power imperatives. Molex PowerPlane connectors are engineered to deliver reliable and efficient power delivery, supporting the seamless integration of diverse server form factors and improving data center operations.

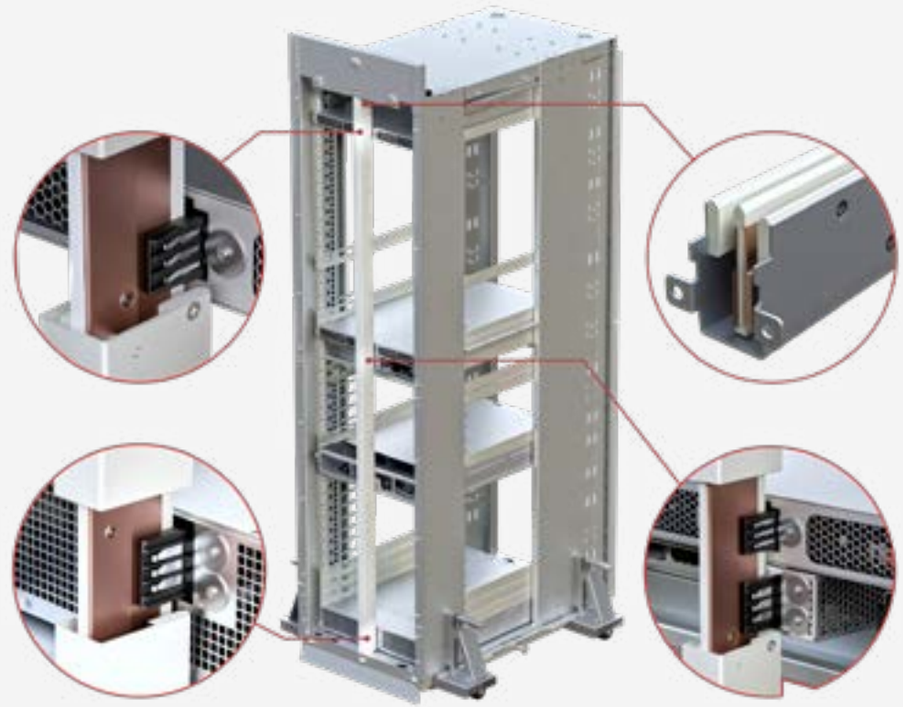

	Current	Max. Voltage Drop	Supported Panel Thickness
OCP ORV3 Power Shelf Molex Output Connector Power Shelf (500A+)	500A+	20mV at 500A	1.10 to 1.32 mm
OCP ORV3 Molex IT Gear Input Connector (100A+)	100A+	50mV at 100A	0.90 to 1.32 mm

MOLEX: YOUR PARTNER IN DATA CENTER TRANSFORMATION

An interconnect provider's commitment to alignment with OCP standards is not just beneficial, but essential to ensure their customers stay ahead in today's data center landscape. By actively contributing to OCP subprojects such as DC-MHS, OAI and ORV3, Molex plays an important role in defining the future of data center infrastructure, and acts as an invaluable co-collaborator to those companies who need to ensure success.

Through our cutting-edge interconnect solutions, Molex empowers data center operators to build world-class, future-proof infrastructure. Through continuous collaboration within the OCP community, the company is dedicated to advancing data center technology and addressing the evolving needs of our customers.

Discover more about how Molex is leading innovation within the OCP ecosystem: **https://www.molex.com/ocp**

molex