DELL EMC

PowerEdge Product Group

Direct from Development

PowerEdge MX7000 Chassis Thermal Airflow Architecture

Tech Note by:
Rick Eiland

SUMMARY
Creating a modular infrastructure that can efficiently cool the high-power, high-density workloads of today and tomorrow requires intelligent, scalable design.

Dell EMC's PowerEdge MX7000 modular infrastructure meets these thermal challenges through an innovative, patent-pending chassis architecture design.

The Multiple Airflow Zone architecture of the new MX7000 chassis enables enhanced energy efficiency in the cooling of higher power and denser system configurations, extends the lifecycle of the chassis to accommodate multiple generations of future IT technology, and delivers excellent investment protection.

Modular infrastructures enable customers to be flexible in configuration, agile in management, and efficient in design, over several generations of compute, storage, and networking needs. The unique thermal design of the new Dell EMC PowerEdge MX7000 chassis ensures that customers can confidently grow over many generations of IT technology, satisfying forthcoming demands for scalable expansion and performance while providing excellent investment protection. The multiple airflow zone architecture of the MX7000 is engineered to ensure fresh air is delivered to each component of the chassis, enabling enhanced cooling of higher power and dense system configurations.

CHASSIS ARCHITECTURE

The overall chassis architecture of the PowerEdge MX7000 uses dedicated, independent airflow paths for each critical subsystem (Compute, I/O, and Power Supplies) to provide fresh air to each individual zone. This design allows for expanded feature support, improved cooling efficiency, and the flexibility to scale with future needs. Figure 1 below provides a view of the front and rear of the MX7000 chassis with each zone highlighted.

- Zone 1 (Blue): Cooling air for the eight vertically-arranged Compute and Storage sled slots at the front of the chassis is drawn thru the chassis by five horizontally-arranged fans at the rear
- Zone 2 (Yellow): Four vertically arranged fans at the front of the chassis push cooling air into the I/O Modules (IOMs) at the rear of the chassis
- Zone 3 (Green): The bottom of the chassis is populated with up to six power supply units (PSUs) each with a dedicated fan and exhaust airflow path out the rear of the chassis

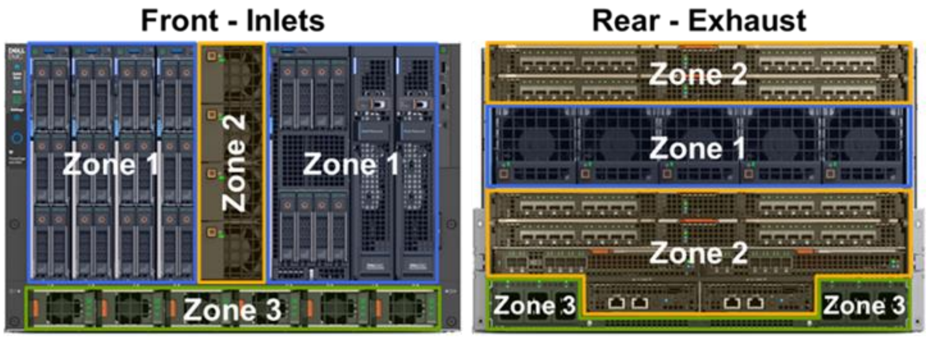

Figure 1: Front and rear view of PowerEdge MX7000 chassis with dedicated airflow zones highlighted

The orthogonal layout of these zones allows for simple, front-to-back airflow thru the chassis. By avoiding overly complex airflow paths to each subsystem, the overall chassis minimizes its impedance, allowing for increased airflow capability compared to traditional modular chassis architectures. Additionally, the independent airflow zones allow for more granular fan control algorithms which increase fan speeds only when and where it is needed to further increase efficiency.

COMPUTE & STORAGE SLED COOLING

Compared with previous modular chassis from Dell and competitors, the MX7000 chassis contains no vertical midplane that would restrict the airflow through the chassis. Instead, sleds within the chassis mate directly with rear IOMs in A1/A2/B1/B2 slots through direct orthogonal interconnects. In the space between the A1/A2 and B1/B2 IOMs, the MX7000 80mm fan modules are directly ducted to the sleds to pull air through compute and/or storage sleds in Zone 1. Multiple chassis seals ensure that the low pressure generated by the rear 80mm fans stays within Zone 1. Containing the low pressure is critical in enabling the increased sled storage density with support for up to six 2.5" HDDs in the front end of the Dell EMC PowerEdge MX740c.

Figure 2 below provides a cross-sectional view of the chassis to highlight the Zone 1 airflow path thru a compute sled. Cool air enters thru the hard disk drive (HDD) bay of the sled before cooling the CPUs, memory, and peripheral components. The 80mm fan modules at the rear of the chassis pull air through the sleds and exhaust hot air out of the system. Fans are positioned in line with the critical heat loads, the CPUs, to avoid complex ducting and provide efficient cooling. Since there are no downstream components in this flow path, a large temperature difference is allowed across CPUs to maximize processor TDP support with headroom to scale with future generations.

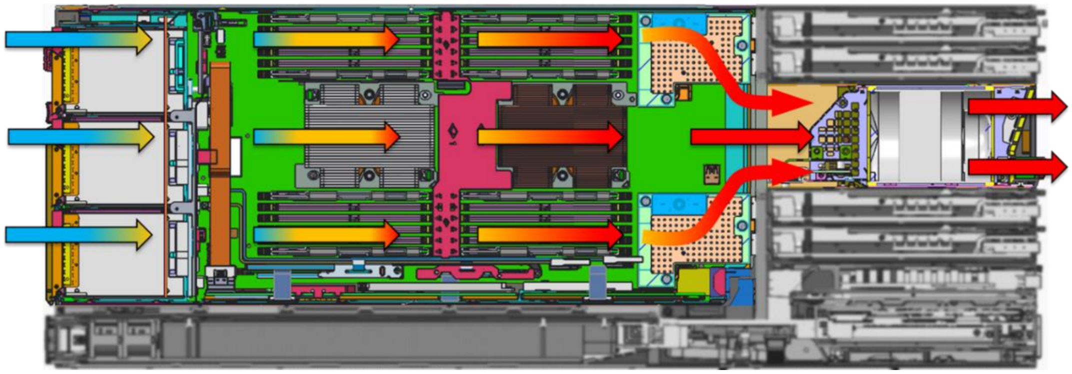

Figure 2: Cross-sectional view of airflow through Zone 1 of the MX7000 chassis, showing cooling of Compute sleds. Air enters at front of chassis (on left) and exhausts thru rear of chassis (on right). Other airflow zones have been greyed out in this graphic.

NETWORKING AND I/O

Due to their internal layout, modular infrastructures have traditionally been forced to cool networking fabrics with preheated air from upstream components or through tortuous pathways that limit the amount of cooling airflow available. With limited airflow capability, the ability to support future networking and I/O technologies becomes more difficult as these technologies continue to advance at a rapid pace. In contrast, the MX7000 chassis has been designed for fresh inlet air delivery to all rear IOMs and Chassis Management Modules (CMMs) via four dedicated fan modules in the front of the chassis. The direct airflow path is shown in the cross-sectional graphic in Figure 3 (at the top of Page 3), which highlights airflow Zone 2. The four 60mm fan modules arranged vertically share a large plenum down the center of the chassis. An air ducts splits the airflow into an upper path which cools IOMs A1/A2 and a lower path to cool IOMs B1/B2/C1/C2 and CMMs. The fresh air entering these modules ensures the MX7000 chassis can run fans at lower speeds and operate over a wider range of environmental conditions as well as providing sufficient airflow capability for future networking technologies.

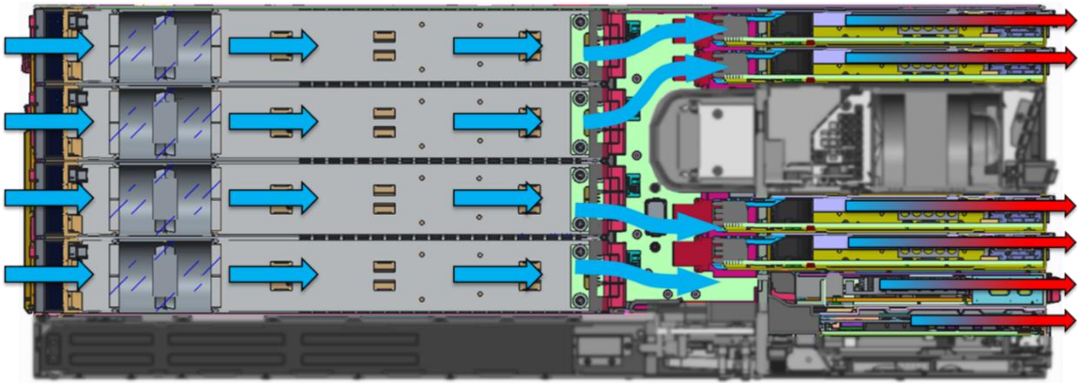

Figure 3: Airflow through Zone 2 of the MX7000 chassis, showing cooling of IOMs. Other airflow zones have been greyed out in this graphic.

POWER SUPPLIES

Power supplies in the MX7000 chassis benefit from a simple, independent airflow path at the bottom of the chassis, as shown in the cross sectional view of Figure 4 below. PSUs intake fresh air at the front of the chassis. Dedicated fans in the front of each PSU push airflow through dense electrical components and under the power connector at the back of the PSU. Since there are no downstream components to cool, PSUs can operate with very high exhaust temperatures of 60°C or higher. This thermal design helps ensure reliable delivery of up to 3000W supplied power in the dense form factor of each MX7000 PSU.

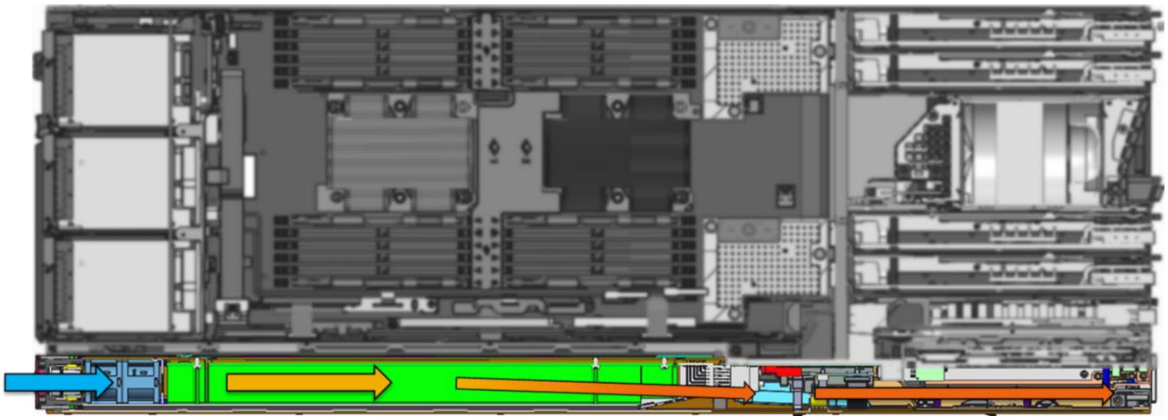

Figure 4: Airflow through Zone 3 of the MX7000 chassis, illustrating cooling of PSUs. Other airflow zones have been greyed out in this graphic.

CONCLUSIONS

Traditional modular architectures have been characterized by complex airflow pathways that restrict airflow available to cool internal components, or require downstream components to operate within the constraints of preheated air from upstream components. This results in sub-optimal energy efficiency as well as places limitations on feature support, which in turn adversely impacts scalable growth for higher performance and/or capacity expansion.

The innovative multiple airflow zone architecture of the new PowerEdge MX7000 chassis overcomes these challenges and ensures that fresh air is delivered to each component in the chassis, enhancing cooling for higher power and denser system configurations. This design enables increased energy efficiency, extends the lifecycle of the chassis to accommodate multiple generations of future IT technology, and delivers excellent investment protection.