DELL Technologies
Server and Infrastructure Engineering

Direct from Development

Multi Vector Cooling 2.0 for Next-Generation PowerEdge Servers

Tech Note by

Matt Ogle
Hasnain Shabbir

Summary

Next-generation PowerEdge servers (15G) support the latest compute, storage and networking technologies with the help of innovation in hardware and thermal controls design that builds on the foundations of the previous-generation (14G) MVC 1.0 solution.

This DfD outlines the new MVC 2.0 innovations on both the hardware thermal design and system thermal controls front that enables maximum system performance, with an eye on thermal efficiency and key customizations desired by customers to tune the system to their deployment needs and challenges.

Introduction

Next-generation PowerEdge servers (15G) support higher-performance CPUs, DIMMs and networking components that will greatly increase the servers' capabilities. However, as capabilities increase, so does the need for continued innovation to keep the system cool and running efficiently.

Multi Vector Cooling (MVC) is not any specific feature – rather it is a term that captures all of the thermal innovations implemented onto PowerEdge platforms. MVC 2.0 for next-generation PowerEdge servers builds upon existing innovations with additional support in hardware design, improved system layout, and cutting-edge thermal controls. These improvements address the needs of an ever-changing compute landscape, demanding a 'green performance', low carbon footprint, as well adding customization levers to optimize not only at the server level, but also at the data center level, generally with airflow handling and power delivery.

Hard Working Hardware

While most of the innovations for MVC 2.0 center around optimizing thermal controls and management, the advancement of physical cooling hardware and its architecture layout is clearly essential:

- **Fans -** In addition to the cost-effective standard fans, multiple tiers of high performing, Dell-designed fans are supported to increase system cooling. The high performance silver and gold fans can be configured into next-generation PowerEdge servers for supporting increased compute density. Figure 1 below depicts the airflow increase (in CFM) for these high performance fans when compared to baseline fans.

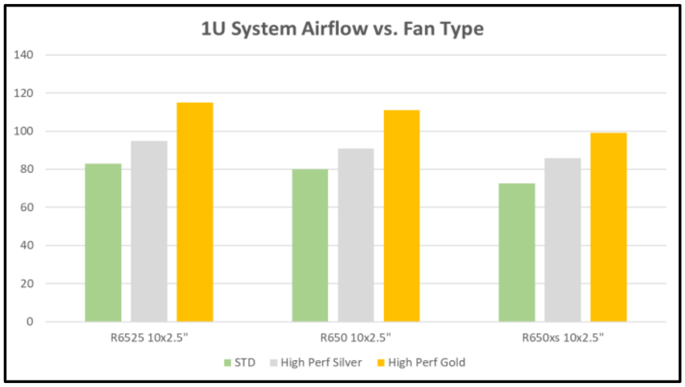

Figure 1 – Comparison of airflow output in CFM

DELL Technologies

- **Heatsinks -** The improved CPU heatsink design not only improves CPU cooling capability, but also helps in streamlining airflow and air temperature distribution across the chassis. Innovative heatsink 'arms' with high performance heat pipes and optimized fin spacing achieve this goal.
- **Layout -** The T-shape system motherboard layout, along with PSUs that are now located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heatsinks, and the streamlined airflow helps with PCIe cooling as well enabling support for PCIe Gen4 adapters. Lastly, this layout creates a better cable routing experience on the PDU side of the racks where power cables are generally separated by grid assignments for redundancy.

AI Based Thermal Controls

To best supplement the improved cooling hardware, the PowerEdge engineering team focused on developing a more autonomous environment. Key features from prior-generations were expanded upon to deliver thermal autonomous solutions capable of cooling next-generation PowerEdge servers.

Our AI based proprietary and patented fuzzy logic driven adaptive closed loop controller has been expanded to not just do fan speed control based on thermal sensor input but is now utilized for power management. This allows for the optimization of system performance, especially in transient workloads and systems operating in challenging thermal environments by automating power management that is required beyond fan speed control for thermal management.

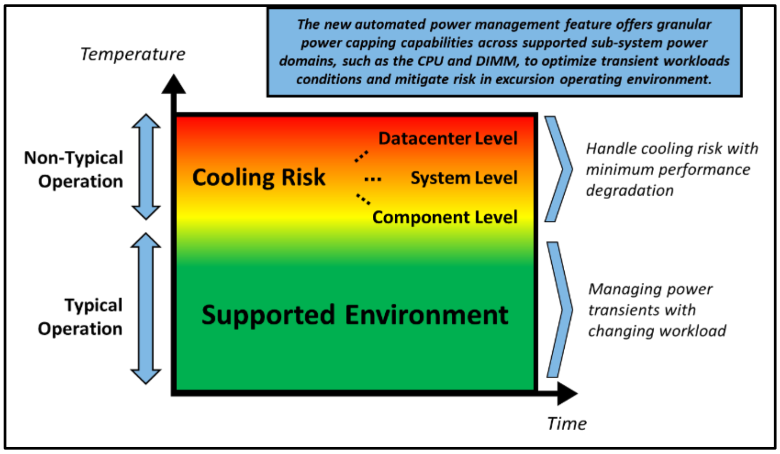

Figure 2 – Each operating environment has unique challenges

This automation with granular power capping capability across various supported sub-system power domains (more specifically CPU and DIMM) ensures thermal compliance with minimum performance impact in challenging thermal conditions. See Figure 2 for illustrates area where new controls solution optimize system performance and uptime.

iDRAC Datacenter Thermal Management Features and OME

With introduction of iDRAC Datacenter license and OME's power manager one-to-many capabilities, customers can monitor and tackle challenges associated to server customizations as well as deployment in their datacenter (power and airflow centric). Below list highlights some of the key features:

1. **System Airflow Consumption -** Users can view real-time system airflow consumption (in CFM), allowing airflow balancing at the rack and datacenter level with newly added integration in the OME Power Manager
2. **Custom Delta-T -** Users can limit the air temperature rise from the inlet to exhaust to right-size their infrastructure level cooling
3. **Custom PCIe inlet temperature -** Users can choose the right input inlet temperature to match 3rd party device requirements

4. **Exhaust Temperature Control -** Users can specify the temperature limit of the air exiting the server to match their datacenter hot aisle needs or limitations (personnel presence, networking/power hardware)
5. **PCIe airflow settings -** Users are provided a comprehensive PCIe devices cooling view of the server that informs and allows cooling customization of 3rd party cards

Figure 3 illustrates how the features previously mentioned work together at a system level:

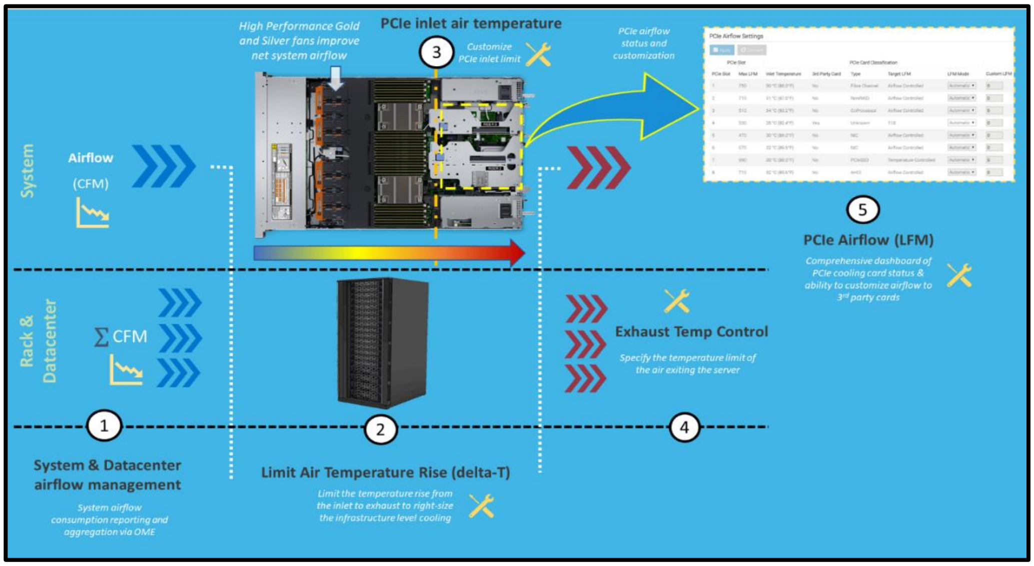

Figure 3 – iDRAC thermal management features and customizations

Channel Card Support

Dell Technologies also offers flexibility for customers wanting to implement non-Dell channel cards. Comprehensive support for PCIe communication standards like PLDM, NC-SI and custom implementations by vendors for GPUs and accelerators, such as Nvidia, AMD, Intel for temperature monitoring and closed loop system fan control. Channel cards that follow these standards will therefore have optimal thermal and power management behavior in PE Servers. Future updates would also include support of new open loop cooling levels defined in latest release of PCIe-SIG standards document.

Conclusion

The Dell Technologies MVC 2.0 solution enables next-generation (15G) PowerEdge servers to support dense configs and workloads with higher-performance cooling hardware, increased automation, simplified but advanced management and channel card flexibility. By expanding upon the existing MVC 1.0 design strategy, the MVC 2.0 solution resolves new thermal challenges so that PowerEdge customers can fully utilize their datacenters while managing the deployment constraints like airflow and power delivery in an optimal fashion.

PowerEdge DfD Repository For more technical learning	**Contact Us** For feedback and requests	**Follow Us** For PowerEdge news

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved.
Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries