



PowerEdge Multi-Vector Cooling

Tech Note by:

Hasnain Shabbir

Tad Walsh

SUMMARY

Higher performance components and denser configurations create challenges for server cooling. PowerEdge Multi-Vector Cooling brings a new approach to server thermal control.

Custom-designed heat sinks and fans, combined with sophisticated system layout and fan zone mapping, increase airflow and cooling through the chassis.

A Dell-developed, intelligent and adaptive, closed-loop algorithm, new with 14G servers, optimizes fan operation in response to a range of parameters including thermal state, fan power, airflow consumption and acoustics.

Moreover, a new PCIe airflow customization capability enables users to specify delivery of airflow to each PCIe slot.

As servers become increasingly powerful and dense, cooling them becomes increasingly challenging as well. Server support for higher performance processors and memory, combined with larger numbers of internal storage drives and higher power PCIe adapters, requires innovative engineering to keep the server cool. PowerEdge Multi-Vector Cooling, designed into PowerEdge 14th-generation (14G) servers, addresses and resolves the full range of issues associated with server thermal control and cooling.

Multi-faceted approach to thermal control

Gone are the days when server cooling could be accomplished by simply running fans faster to move more air *en masse* through the chassis. Servers must be capable of moving air variably over internal components, providing optimum cooling when necessary to keep each component and the server overall within operating parameters. Moreover, user requirements now also include parameters for power consumption, airflow consumption and acoustics. PowerEdge Multi-Vector Cooling resolves these issues with a multi-faceted approach to thermal control.

On one hand, PowerEdge Engineering has developed custom-designed heat sinks and fans, while at the same time advancing internal component layout to improve airflow and cooling:

- Highly reliable fans, customized to Dell standards for efficiency, performance, long life, vibration, and acoustics. PowerEdge fans are capable of long life (beyond 5 years) even if the fans run at full speed all the time.
- Dell-designed custom heatsinks optimize component cooling at minimum required airflow, and yet support high performance.
- The system layout of PowerEdge servers is designed to increase airflow through the chassis, enabling enhanced cooling of higher power and dense system configurations.

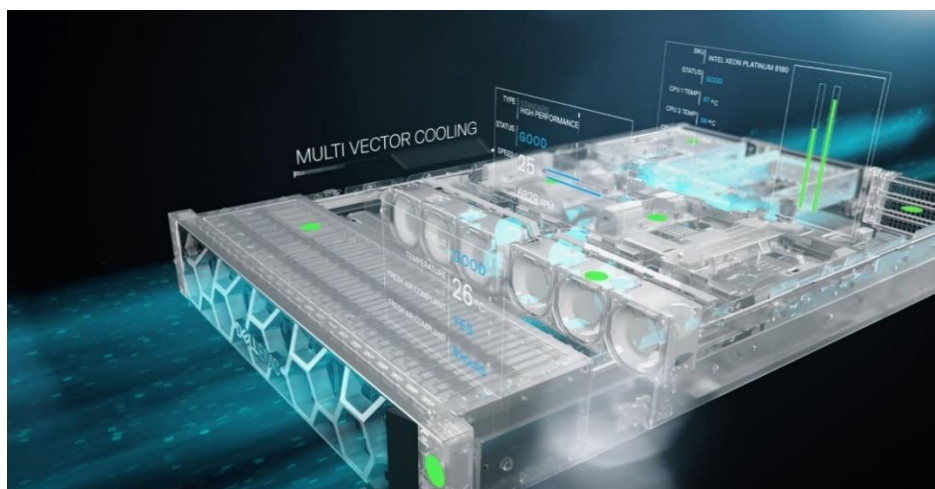


Figure 1: PowerEdge Multi-Vector Cooling – It's not just about temperature



This latter factor, advanced thermal design, streamlines the airflow pathways within the server, directing the appropriate volume of air to where it is needed inside the chassis. This is accomplished in part by a large set of sensors that detect and allow accurate interpretation of real-time system thermal state at various locations within the server. Through Dell-patented *fan zone mapping* (based on location of e.g. CPU's, memory, mini-PERC, on-board sensors, network daughter cards, etc.), components can be targeted for cooling, resulting in maximum performance of the component with optimal airflow and power efficiency.

At the same time, an intelligent and adaptive closed-loop control algorithm, also new with PowerEdge 14G servers, optimizes fan operation in response not only to the thermal state detected by the sensors just mentioned, but also to factors including fan power, airflow consumption and acoustics, all while maintaining component temperatures within targeted parameters. This Dell-developed algorithm (patents pending) is based on fuzzy logic, and responds to inputs from the large number of sensors mentioned above.

PCIe Airflow Settings

In addition, given the increasing power consumption and heat produced by advanced PCIe adapters becoming available in the marketplace, PowerEdge Engineering has developed an innovative solution for PCIe adapter cooling that gives users an accurate picture of airflow delivered to *each* PCIe slot, as well as the capability to customize the delivery of airflow for their custom PCIe cards.

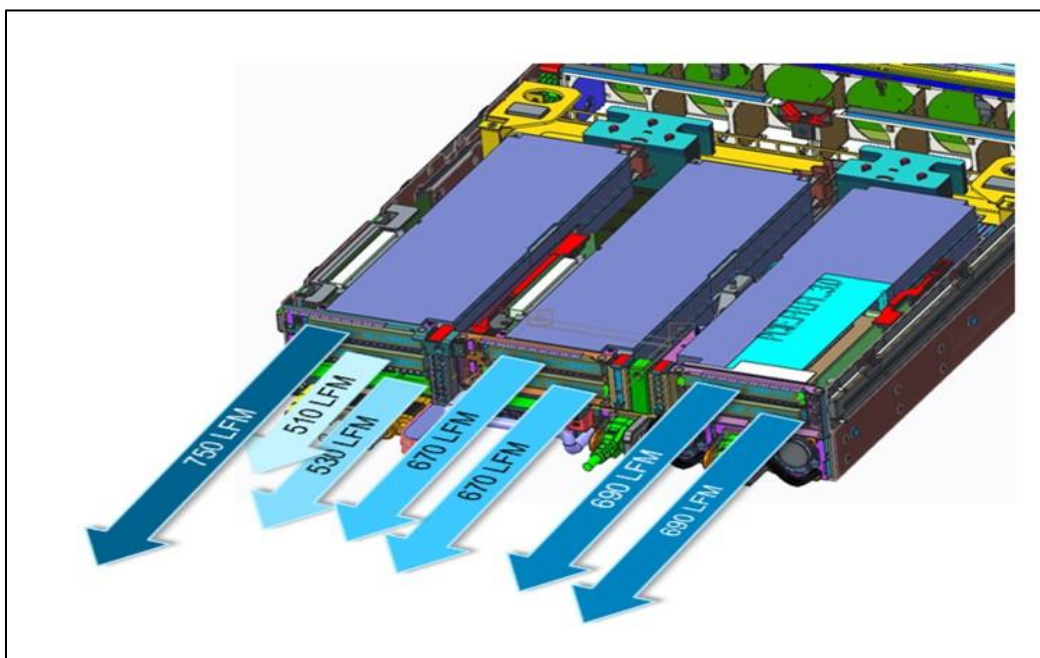


Figure 2: PowerEdge 14G servers enable users to customize delivery of airflow to each PCIe slot.

PCIe airflow settings deliver an accurate representation of slot-by-slot PCIe airflow. This information is reported in terms of Linear Feet per Minute (LFM), a generally accepted industry standard metric on how PCIe card airflow is specified. The LFM metrics are displayed in various iDRAC interfaces, allowing users to be aware of:

- the maximum LFM capability of each slot within the server
- what approach is being taken for PCIe cooling for each slot (e.g. airflow controlled; temperature controlled)
- and, if the installed adapter is a 3rd Party Card (customer custom card), the minimum LFM being delivered to the slot.



With all this information, the user can access the PCIe airflow customization capability in the iDRAC web GUI (see Figure 3 below) and specify a custom minimum LFM value for the 3rd Party adapter, allowing more accurate definition of the adapter's cooling needs. Users can also cross-reference this value with the specifications provided by the adapter vendor (e.g. in the Spec Sheet for the adapter): The LFM airflow value for the 3rd Party adapter provided by the algorithm may be more or less than the cooling needs specified by the vendor. In this case, the user can simply fine-tune the value entered for the adapter by entering the LFM provided by the vendor.

Apply

Discard

PCIe Slot		PCIe Card Classification				
PCIe Slot	Max LFM	3rd Party Card	Type	Target LFM	LFM Mode	Custom LFM
1	750	No	PCIeExtender	Airflow Controlled	Automatic	0
3	510	Yes	Fiber Channel	100	Custom	100
4	530	No	Video	Temperature Controlled	Automatic	0
6	670	No	RAID	Temperature Controlled	Automatic	0
7	690	Yes	NonRAID	200	Automatic	0
8	710	No	PCIeSSD	Airflow Controlled	Automatic	0

Figure 3: Users can customize airflow to each PCIe slot using the PCIe Airflow Settings screen in iDRAC

Additional capabilities within the server cooling customization options include:

- Custom thermal settings (e.g. Thermal Profiles such as Maximum Performance; Maximum Performance per Watt; or Sound Cap)
- Custom fan speed options (minimum fan speed; fan speed offsets)
- Custom exhaust temperature setting

The iDRAC web GUI or RACADM also displays a real-time system airflow metric (in cubic feet per minute, of CFM) to the user to enable datacenter airflow balancing based on aggregation of per server CFM.

Sound Cap – where acoustics is a key requirement

Another new feature of PowerEdge 14G servers, associated with fan operation and Multi-Vector Cooling, is Sound Cap. Sound Cap was developed in response to customer requests, and is for specialized environments in which minimizing acoustical output is a higher criteria than peak raw performance. Sound Cap limits, or “caps”, CPU power consumption and thus fan speed and acoustical ceiling. Its application is unique for acoustical deployments and may result in reduced system performance.

Conclusion

Building on the wealth of knowledge accumulated over years of server design, PowerEdge Multi-Vector Cooling is a step-function improvement in capability above previous-generation servers. New in PowerEdge 14G servers, Multi-Vector Cooling enables deeper and more granular resolution of customer pain points concerning PCIe cooling, airflow consumption, power consumption, and acoustics, helping to ensure the high standards of reliability, performance, and efficiency expected of a PowerEdge.