

The Future of Server Cooling

Part 1. The History of Server and Data Center Cooling Technologies

Summary

Today's servers require more power than ever before. While this spike in power has led to more capable servers, it has coincidentally pushed legacy thermal hardware to its limit. The inability to support top-performance servers without liquid cooling will soon become an industry-wide challenge. We hope that by preparing our PowerEdge customers for this transition ahead of time, and explaining in detail why and when liquid cooling is necessary, they can easily adapt and get excited for the performance gains liquid cooling will enable.

Part 1 of this three part series, titled *The Future of Server Cooling*, covers the history of server and data center thermal technologies - which cooling methods are most commonly used, and how they evolved to enable the industry growth seen today.

Authors

Matt Ogle
Todd Mottershead

The Future of Server Cooling was written because the next-generation of PowerEdge servers (and succeeding generations) may require liquid cooling assistance to enable certain (dense) configurations. Our intent is to educate customers about why the transition to liquid cooling is inevitable, and to prepare them for these changes. Integrating liquid cooling solutions on future PowerEdge servers will allow for significant performance gains from new technologies, such as next-generation Intel® Xeon® and AMD EPYC™ CPUs, DDR5 memory, PCIe Gen5, DPUs, and more.

Part 1 of this three part series reviews some major historical cooling milestones to help explain why change has always been necessary. It also describes which technologies have evolved over time to advance to where they are today - the historical evolution of thermal technologies for both the server and the data center.

Data centers cannot exist without sufficient cooling

A data center comprises many individual pieces of technology equipment that work together collectively to support continuously running servers within a functional facility. Most of this equipment requires power to operate, which converts electrical energy into heat energy as it is used. If the heat generated grows too large, it can create undesirable thermal conditions, which can cause component and server shutdown, or even failure, if not managed properly.

Cooling technologies are implemented to manage heat build-up by moving heat away from the source (because heat cannot magically be erased) and towards a location where it is safely dispersed. This allows technology equipment within the data center to continue to work reliably and uninterrupted from the threat of shutdown from overheating. Servers from Dell Technologies can automatically adjust power consumption, but without an effective cooling solution the heat buildup within the data center would eventually exceed the capability of the server to operate, creating enormous financial losses for business.

Two areas of coverage

Cooling technologies are typically designated to two areas of coverage - directly inside of the server and at the data center floor. Most modern data centers strategically use cooling for both areas of coverage in unison.

- Cooling technologies located directly inside of the server focus on moving heat away from dense electronics that generate the bulk of it, including components such as CPUs, GPUs, memory, and more.
- Cooling technologies located at the data center floor focus on keeping the ambient room temperature cool. This ensures that the air being circulated around and within the servers is colder than the hot air they are generating, effectively cooling the racks and servers through convection.

Legacy Server Cooling Techniques

Four approaches have built upon each other over time to cool the inside of a server: conduction, convection, layout, and automation, in chronological order. Despite the advancements made to these approaches over time, the increasing thermal design power (TDP) requirements have made it commonplace to see them all working together in unison.

Conduction was the first step in server cooling evolution that allowed the earliest servers to run without overheating. Conduction directly transfers heat through surface contact. Historically, conduction cooling technologies, such as heat spreaders and heat sinks, have moved heat away from server hot spots and stored it in isolated regions where it can either reside permanently, or be transferred outside of the box through an air or liquid medium. Because heat spreaders have limited capabilities, they were rapidly replaced by heat sinks, which are the industry standard today. The most effective heat sinks are mounted directly to heat producing components with a flush base plate. As development advanced, fins of varying designs (each having unique value) were soldered to the base plate to maximize the surface area available. The base plate manufacturing process has shifted from extrusion to machine or die-cast, which reduces production time and wasted material. Material changed from solely aluminum to include copper for use cases that require its ~40% higher thermal conductivity. The following figure provides an example:




Figure 1. Heat sink base plate uses copper to support higher power

Convection cooling technologies were introduced to server architecture when conduction cooling methods could no longer solely support growing power loads. Convection transfers heat outside of the server through a medium, such as air or liquid. Convection is more efficient than conduction. When the two are used together, they form an effective system - conduction stores heat in a remote location and then convection pushes that heat out of the server.

Technologies such as fans and heat pipes are commonly used in this process. The evolution of fan technology has been extraordinary. Through significant research and development, fan manufacturers have optimized the fan depth, blade radius, blade design, and material, to present dozens of offerings for unique use cases. Factors such as the required airflow (CFM) and power/acoustic/space/cost constraints then point designers to the most appropriate fan. Variable speed fans were also introduced to control fan speeds based on internal temperatures, thereby reducing power usage. Heat pipes have also undergone various design changes to optimize efficiency. The most popular type has a copper enclosure,

sintered copper wick, and cooling fluid. Today they are often embedded in the CPU heatsink base, making direct contact with the CPU, and routing that collected heat to the top of the fins in a remote heatsink.

Layout refers to the placement and positioning of the components within the server. As component power requirements increased at a faster rate than conduction and convection technologies were advancing, mechanical architects were pressed to innovate new system layout designs that would maximize the efficiency of existing cooling technologies. Some key tenets about layout design optimization have evolved over time:

- Removing obstructions in the airflow pathway
- Forming airflow channels to target heat generating components
- Balancing airflow within the server by arranging the system layout in a symmetrical fashion

Automation is a newer software approach used to enable a finer control over the server infrastructure. An autonomous infrastructure ensures that both the server components and cooling technologies are working only as hard as needed, based on workload requirements. This lowers power usage, which reduces heat output, and ultimately optimizes the intensity of surrounding cooling technologies. As previously mentioned, variable fan speeds were a cornerstone for this movement, and have been followed by some interesting innovations. Adaptive closed loop controllers have evolved to control fan speed based on thermal sensor inputs and power management inputs. Power capping capabilities ensure thermal compliance with minimum performance impact in challenging thermal conditions. For Dell PowerEdge servers, iDRAC enables users to remotely monitor and tackle thermal challenges with built-in features such as system airflow consumption, custom delta-T, custom PCIe inlet temperature, exhaust temperature control, and adjustment of PCIe airflow settings. The following figure illustrates the flow of these iDRAC automations:

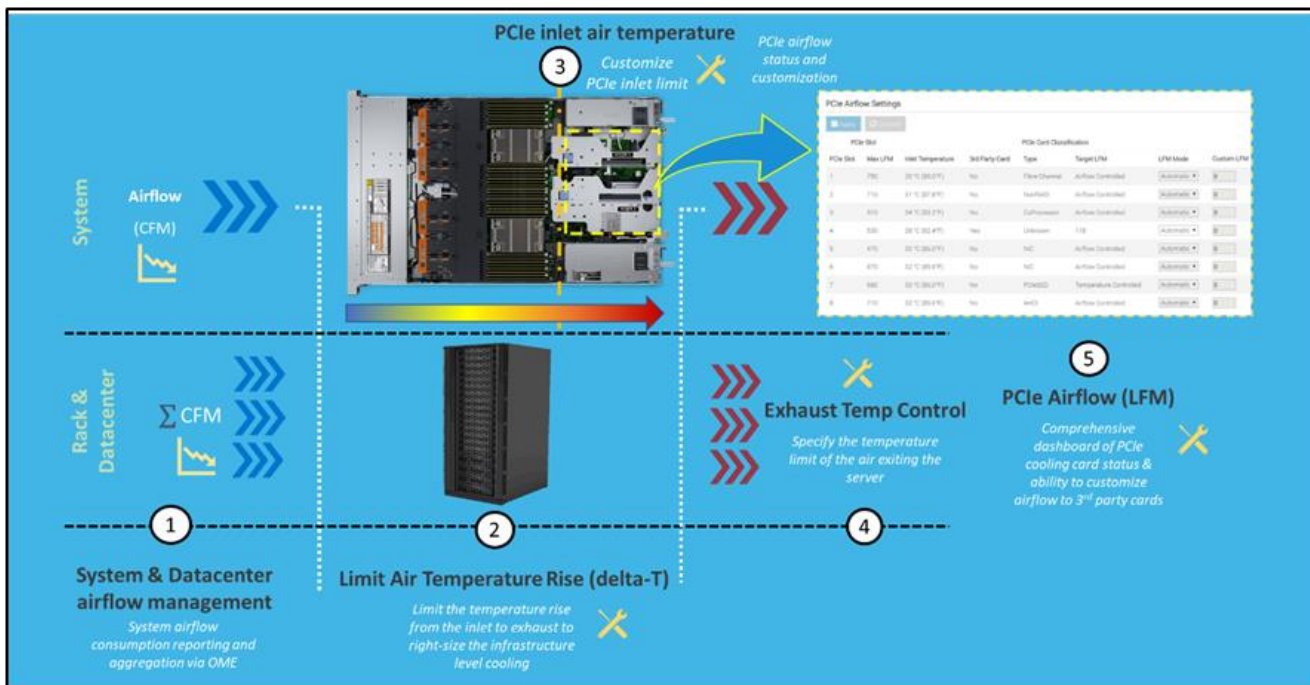


Figure 2. Thermal automations enabled by Dell proprietary iDRAC systems management

Legacy data center cooling techniques

Heat transfer through convection is rendered useless if the intake air being moved by fans is not colder than the heated air within the server. For this reason, cooling the data center room is as important as cooling the server: the two methods

depend on one another. Three main approaches to data center cooling have evolved over time – raised floors, hot and cold aisles, and containment, in chronological order.

Raised floors were the first approach to cooling the data center. At the very beginning, chillers, and computer room air conditioning (CRAC) units were used to push large volumes of cooled air into the datacenter, and that was enough. However, the air distribution was wildly unorganized and chaotic, having no dedicated paths for hot or cold airflow, causing many inefficiencies such as recirculation and air stratification. Because adjustments were required to accommodate increasing power demands, the data center floor plan was redesigned to have raised floor systems with perforated tiles replacing solid tiles. This provided a secure path for the cold air created by CRAC units to stay chilled as it traveled beneath the floor until being pulled up the rack by server fans.

Hot and cold aisle rack arrangements were then implemented to assist the raised floor system when the demands of increasing heat density and efficiency could not be met. This configuration has cool air intakes and warm air exhausts facing each other at each end of a server row. Convection currents are then generated, which helped to improve airflow. However, this configuration was still unable to meet the demands of growing data center requirements, as airflow above the raised floors remained chaotic. Something else was needed to maximize efficiency.

Containment cooling ideas propagated to resolve the turbulent nature of cool and hot air mixing above raised floors. By using a physical barrier to separate cool server intake air from heated server exhaust air, operators were finally able to maintain tighter control over the airstreams. Several variants of containment exist, such as cold aisle containment and hot aisle containment, but the premise remains the same – to block cold air from mixing with hot air. Containment cooling successfully increased data center cooling efficiency, lowered energy consumption, and even created more flexibility within the data center layout (as opposed to hot and cold aisle rack arrangements, which require the racks to be aligned in a certain position). Containment cooling is commonly used today in conjunction with raised floor systems. The following figure illustrates what a hot aisle containment configuration might look like:




Figure 3. Hot aisle containment enclosure diagram, sourced from Uptime Institute

What's Next?

Clearly the historical evolution of these thermal techniques has aided the progression of server and data center technology, enabling opportunities for innovation and business growth. Our next-generation of PowerEdge servers will see technological capabilities jump at an unprecedented magnitude, and Dell will be prepared to get our customers there with the help of our new and existing liquid cooling technologies. Part 2 of this three part series will discuss why power requirements will be rising so aggressively in our next generation PowerEdge servers, what benefits this will yield, and which liquid cooling solutions Dell will provide to keep our customers' infrastructure cool and safe.



[Learn more](#) about
Dell PowerEdge
servers



[Contact us](#) for
feedback and
requests



Follow us for
PowerEdge
news