

Advanced Thermal Control: Optimizing across Environments and Power Goals

This Dell technical white paper explains how thermal control is used to reduce operating costs and adapt to user environments.

Paul Artman

Chris E. Peterson



Contents

Executive summary.....	3
Introduction	3
Advanced thermal control	4
Supported server environments.....	5
Extended temperature range	5
Support for acoustically sensitive installations.....	5
Thermal settings	7
Thermal base algorithm	8
User options.....	9
Summary	10

This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

February 2012 | Rev 1.0

Executive summary

The thermal control of today's server has become increasingly sophisticated to provide the best experience for your deployment environment. The benefits of this approach are seen in reduced system power and acoustics with a broader feature-set support. This approach has also allowed the supported temperature range to extend to as high as 45°C (113°F). With this increased thermal control, more capability has been provided to help you tailor the thermal control approach to your environment.

Throughout this document, thermal control will be referenced. Thermal control is the active management of system cooling through fan speed and system power management to ensure system reliability while minimizing system power consumption, airflow, and system acoustic output.

Introduction

Deployment of today's server technology continues to push the bounds of what can be supported by many data centers. IT professionals and data center managers are facing challenges in both power delivery and cooling capacity based on server power requirements. In addition, servers are increasingly being deployed in non-data center environments where there is a higher sensitivity to acoustics.

Power and airflow consumed to cool a server can be significant contributors to overall power consumption of the data center. Dell engineers have designed a thermal solution that minimizes these contributions. In addition, Dell has focused attention on system acoustics to allow support of server tower configurations in office spaces where they can be in very close proximity to the user. The current generation of Dell™ PowerEdge™ servers achieves this by incorporating sophisticated thermal control, strategic component placement and isolation, airflow management, and power-efficient fans.

Some of the revolutionary improvements in Dell's thermal design and control include the following:

- The amount of detailed information has increased for server power, temperature, and configuration, resulting in:
 - > Minimized total system power; for example, fan power consumption is balanced against component requirements
 - > Reduced fan speeds and airflow
 - > Lower sound power levels from previous server generations
- Advanced power management has been used to expand supported ambient temperature with minimum impact to reliability and performance.
- Improved tools have been developed to tailor the system thermal response to your needs, including:
 - > Ability to adjust thermal control setting and to optimize against system performance and performance-per-watt needs
 - > Ability to increase fan speeds to reduce the exhaust temperature and to provide increased airflow to high-power PCI cards

Advanced thermal control

Dell has continued to evolve the thermal control for the current portfolio of products. To more accurately meet system and component cooling requirements, discrete real-time component feedback into thermal control has been significantly increased.

Temperature monitoring has been added for hard disk drives (HDD), PowerEdge Expandable RAID Controller (PERC) batteries, the PERC RAID on chip (ROC), graphics processing units (GPU), and Select Network Adapters. In addition, many of Dell's add-in I/O cards output their temperatures. Some systems have the ability to monitor, report, and control to a discrete exhaust temperature. The temperatures of the processor, memory, ambient, board, power supply unit, and chipset continue to be monitored as inputs into thermal control.

In addition to temperature monitoring, Dell's thermal control has the ability to monitor and control system power. Component power monitoring is provided to the thermal control for the processor, memory, I/O, and HDD subsystems along with the total system power. With the ability to control system power, Dell ensures component thermal requirements are met across a wide range of configurations and operating environments. In typical environments, component power management will generally not be seen. However, when the system operates outside Dell-supported temperature ranges, system power may be capped to ensure system uptime and "ride through" in case of a data center cooling failure.

Because processor cooling continues to be the primary driver of system thermal control response for stressed operation, Dell's design focuses on minimizing its impact. Dell monitors processor power and temperature as part of thermal control implementation and to ensure compliance to processor thermal specifications.¹ The results are minimized system power consumption and acoustics during stressed operation.

Processor cooling settings are customized for each server with the goal of minimizing total system power. In some systems, large changes can be made in processor temperatures with little fan power. Available on the Dell PowerEdge R620 and R720, "Dell Solution Optimized CPU Cooling" is a setting that exploits this design. In this setting, total system power consumption reduces while the processor temperatures are lowered through a marginal fan power increase. For other systems, large changes in fan power are required for small changes in processor temperatures. For these systems, the processor may be cooled to the new Intel® DTS 2.0 specification, enabled for the Intel Xeon processor E5 family, which allows the processor to operate at higher temperatures and higher power consumption. By allowing the processor to operate slightly below thermal specifications, fan speeds can be reduced while still ensuring system reliability. This results in lower system power consumption and sound power levels. This is the approach used on the PowerEdge M620 and T620.

¹ "CPU Monitoring With DTS/PECI" by Michael Berktold and Tian Tian, Intel Inc., 2010, <http://download.intel.com/design/intarch/papers/322683.pdf>

Supported server environments

Dell has expanded the range of supported environments to respond to new deployment models, from traditional data centers to Fresh Air data centers and remote installations with no environmental control, and from back rooms to retail and office environments. Extended ambient temperature ranges and lower system sound power levels allow servers to operate in diverse and challenging environments.

Extended temperature range

In the current generation of Dell PowerEdge servers, Dell's design supports the extended temperature range that runs from -5°C to 45°C^2 (23°F to 113°F) for certain configurations and systems. This allows for compliance to latest A3 and A4 temperature-class specifications defined by ASHRAE and being adopted by the EU Code of Conduct Best Practices for Data Centres. Dell's support for extended temperature ranges reduces the need for compressor-based cooling for some data center environments, and in some cases, compressor-based cooling may be eliminated.³ In addition, allowing extended temperature ranges can provide for server deployment in spaces with little or no temperature management, such as manufacturing spaces or service bays, and thus deliver greater flexibility.

The ability of Dell servers to support an extended temperature range was developed through long-term testing and analysis of a population of servers. The server population was tested to the equivalent of a seven-year field life in a very challenging environment with a southern European climate, such as in southern Italy or Greece. The long-term testing included high-temperature and high-humidity conditions to simulate summer; low temperature and low humidity to simulate cool dry winter air; and temperature cycling to simulate daily and seasonal temperature changes. By testing servers to the full extent of their usable field life under worst-case conditions, Dell was able to identify possible areas of concern and address these concerns through design improvement.

All fresh-air rated Dell server models, such as the PowerEdge R720, are also subjected to validation testing in custom environmental chambers at the extreme corners of the fresh air temperature and humidity specification. Also, as part of Dell's design for reliability processes, Dell conducted thorough component de-rating analysis on every component in the server to ensure long-term reliability. De-rating involves checking the temperatures and voltages of the components under worst-case environmental and loading conditions to make sure the components stay well within their datasheet ratings for all usage and environmental conditions.

Support for acoustically sensitive installations

If acoustical performance in a server is important to you, Dell's current generation of PowerEdge servers offers you new options. Dell has set new requirements for acoustical performance of specific

² <http://www.dell.com/freshair>

³ "Chiller-less Facilities: They May be Closer Than You Think," by John Fitch, David Moss, and Paul Artman, Dell Inc., 2011.

configurations. Since fan speed is the primary determiner of server noise levels⁴, fans with significantly lower minimum baseline speeds than those in previous generations were incorporated, and were controlled to operate at the lowest fan speeds to meet simultaneously the new acoustical requirements and cooling needs.

In addition, Dell has continued its rigorous mechanical design for vibration and acoustics around HDDs and fans. In order to reduce vibration induced by fan motors as the fans spin, and thus reduce noises such as tones, buzzes, and rattles, Dell isolates fans from other parts in the server. Because HDDs also vibrate as their media rotate, Dell has designed its server chassis to reduce the vibration to and from the HDDs.

The results are that the most common shipping configurations and the lower-end configurations of present generation of PowerEdge servers are about half as loud as those of the previous generation. For example, the lowest configuration⁵ of the tower server T620 is quiet enough for a quiet library. In fact, the present generation tower servers are all quiet enough for desktops in office environments for the most commonly shipping configurations and lower-end configurations. Even the R620 and R720 rack servers, in the most commonly shipping configurations and lower end configurations, are quiet enough for office environments.

Table 1 shows the acoustical reference points and comparison of acoustical output of the latest-generation PowerEdge T620 to the previous generation PowerEdge T610. Both systems are in the lowest configuration.

Table 1. Acoustical reference points and output comparisons

	Value measured at your ears		Equivalent familiar noise experience
	LpA, dBA, re 20 µPa	Loudness, sones	
	90	80	Loud concert
	75	39	Data center, vacuum cleaner, voice must be elevated to be heard
	60	10	Conversation levels
	45	4	Whispering, open office layout, normal living room
PowerEdge T610 (previous-generation)	35	2	Quiet office
PowerEdge T620 (latest-generation)	30	1	Quiet library
	20	0	Recording studio

⁴ “Dell Enterprise Acoustics,” Chris E. Peterson, Dell, Inc., 2011, available at <http://www.dell.com/downloads/global/products/pedge/en/acoustical-education-dell-enterprise-white-paper.pdf>.

⁵ Lowest bin processor, low-capacity DIMM, no PCI card, 1x SATA HDD

Very early in the product design cycle, Dell conducts acoustical and vibration modeling to predict effectiveness of design approaches. In order to verify that the new acoustical requirements per configurations are being met during product development, sound quality and sound power level measurements are conducted in accordance with ISO7779⁶, ECMA-74⁷, ISO9296⁸, and other standards in Dell-calibrated hemi-anechoic test chambers.

Some components incur greater heat load than others and thus require significantly greater cooling and cause higher noise levels when they are installed. Although you would expect higher acoustical values with higher performance parts, the impact of some components on acoustics may not be intuitive. Installation of high-powered expansion cards like PERCs, GPUs, and some specific PCI cards cause increased noise levels. When high-powered processors are exercised, they also require greater cooling, hence greater noise levels. Despite this, the T620 tower server has been specifically designed to maintain open-office-layout acoustics when a GPU is installed and is idling.

As noted previously, HDDs also contribute to acoustics. Their contributions depend strongly on their rotation speed. For example, 15k HDDs are much louder (more than 30% louder) and produce tones with higher prominence ratios than HDDs that spin more slowly. As the quantity of HDDs increases, the loudness also increases. Therefore, the best choice for HDDs for an acoustically sensitive environment is a single, 7200rpm SATA HDD.

Thermal settings

To match your deployment, environmental, and usage requirements, you can set your BIOS system profiles accordingly.⁹ By default, the thermal base algorithm is set to Auto, which maps to your selected system profiles. However, under the iDRAC Settings>Thermals menu, you can select a custom thermal control, independent of BIOS system profiles. Figure 1 shows the menu for iDRAC's thermal settings.

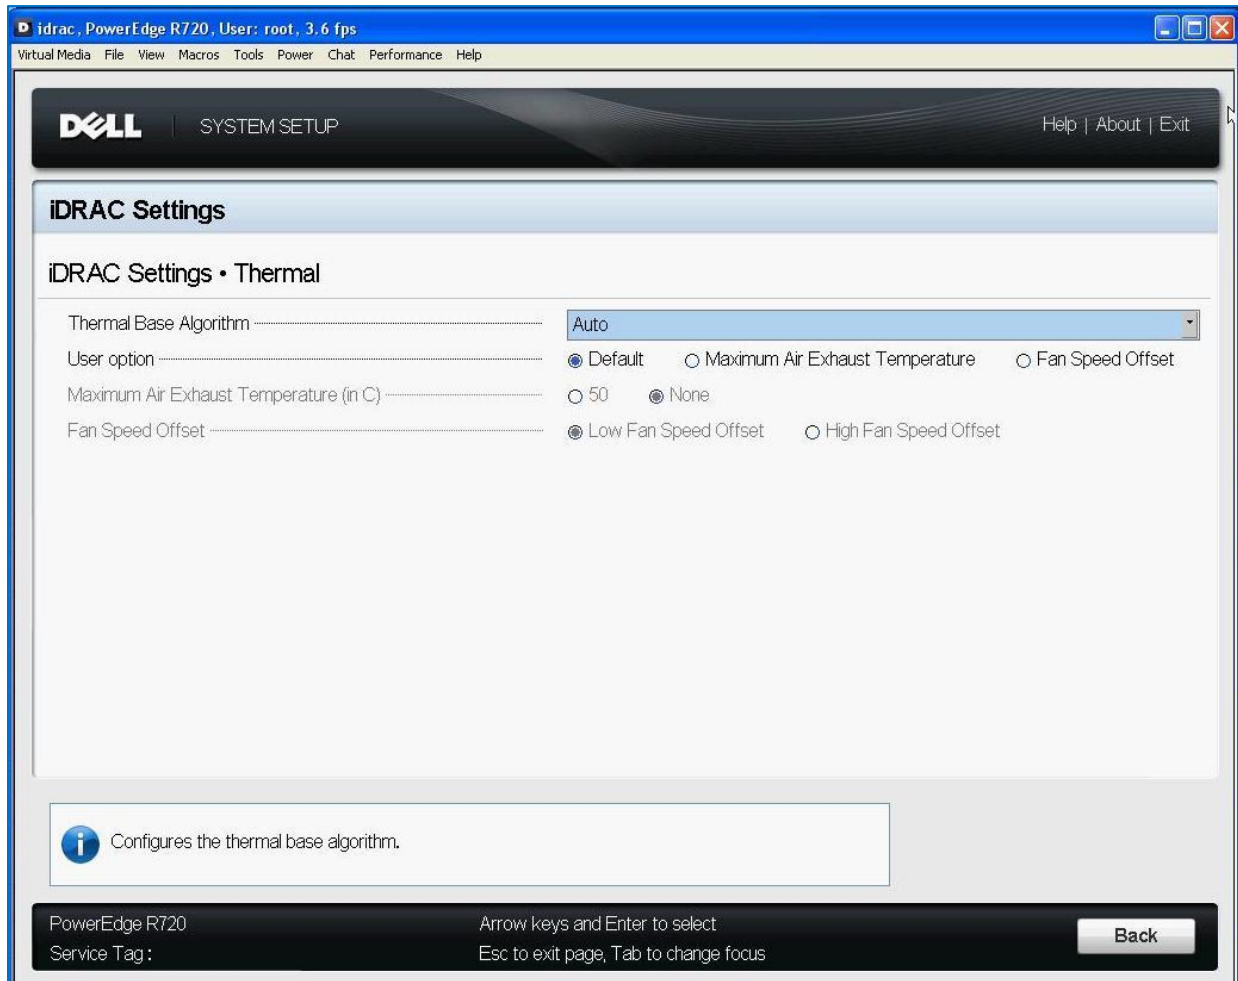
⁶ ISO 7779, 3rd ed., *Acoustics - Measurement of airborne noise emitted by information technology and telecommunications equipment*, Geneva: International Organization for Standardization, 2010.

⁷ ECMA-74, 11th ed., *Measurement of Airborne Noise Emitted by Information Technology and Telecommunications Equipment*, Geneva: ECMA International, 2010.

⁸ ISO 9296:1988 (E) *Acoustics - Declared noise emission values of computer and business equipment*, Geneva: International Organization for Standardization, 1988.

⁹ "Power Efficiency "How To" for 12G PowerEdge Server Portfolio", by John Jenne, Dell Inc., 2012.

Figure 1. iDRAC Settings, Thermal menu

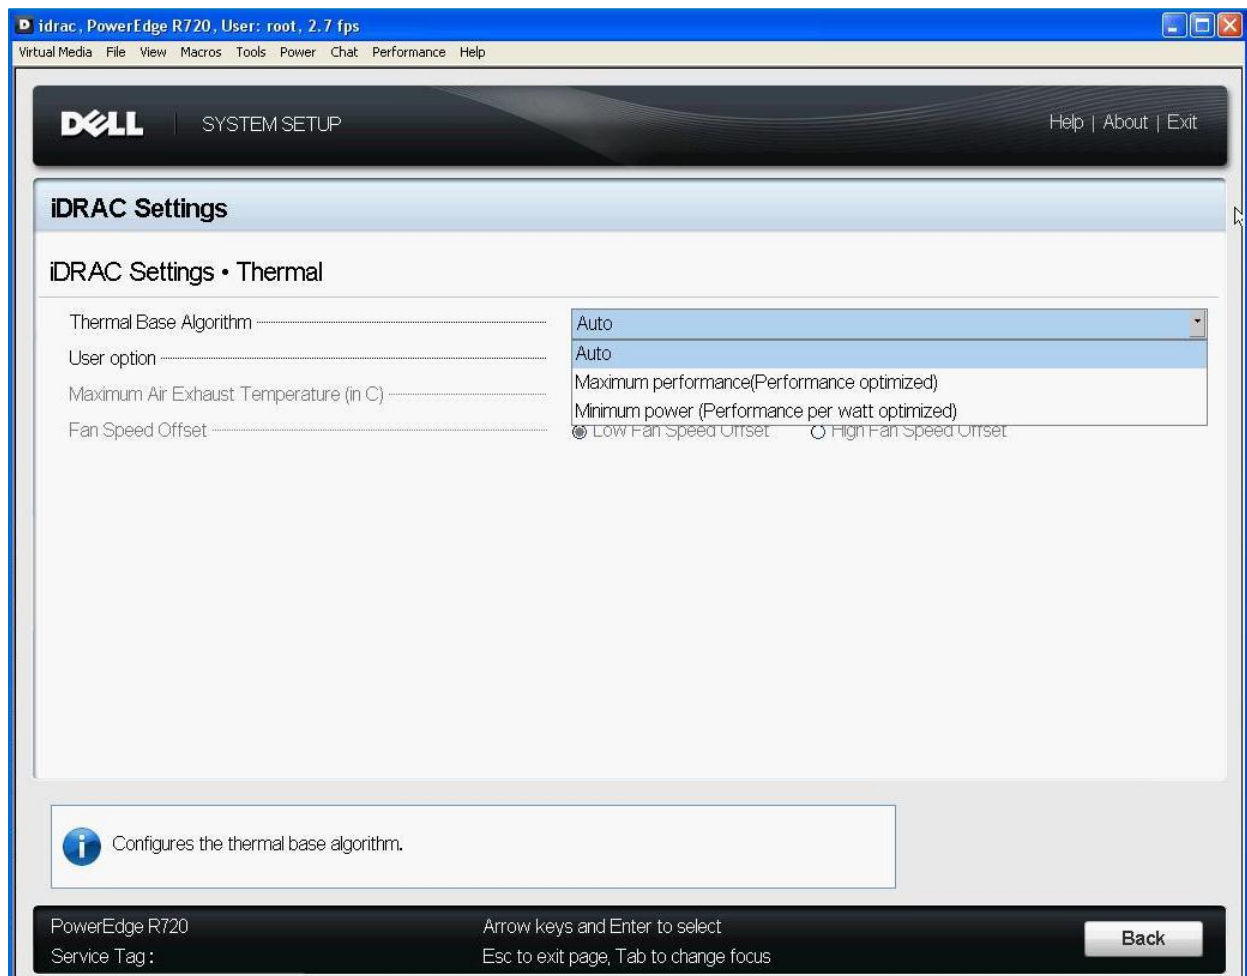


Thermal base algorithm

In the iDRAC Settings, Thermals menu, you can select Auto, Maximum performance (Performance optimized), or Minimum power (Performance per watt optimized). The Auto option defaults to the settings in the BIOS system profile. Figure 2 shows the menu selections for the Thermal Base Algorithm.

The Maximum performance option has more aggressive temperature settings for processor and memory cooling than the Minimum power option. The benefit is the minimization of thermally driven performance impacts at the expense of increased fan power. The Minimum power setting balances component cooling requirements against performance and system power constraints. Thermally driven performance impacts are not expected for either setting for typical data center ambient temperatures (18°C to 30°C; 64.4°F to 86°F). When performance is critical and system operation may occur at elevated temperatures, the Maximum performance setting may provide improved performance.

Figure 2. iDRAC Thermal Base Algorithm settings



User options

Under the iDRAC Setting, Thermal menu, you can select between Default, Maximum Exhaust Temperature, and Fan Speed Offset settings. Only one selection may be made in this menu; each setting is mutually exclusive of the other settings. Of note, the Default setting will have no impact on thermal settings.

Maximum Air Exhaust Temperature setting

The most efficient thermal design often results in system fan speeds at the lowest achievable setting to meet thermal and reliability requirements. This reduces system power and data center airflow consumption by the server and has the added benefit of increasing the temperature difference across the CRAC (computer room air conditioner) coils, which results in higher operating efficiency of the CRAC.

For some environments, this may have impacts on adjacent equipment (for example, the PDU or switch) or the performance. To address these conditions, the Maximum Air Exhaust Temperature setting was developed for R600 class and higher Dell PowerEdge products. This industry-first thermal control setting allows the user to select a maximum exhaust temperature of 50°C (122°F). This feature

uses several discrete exhaust temperature sensors with fan speed control and power management to ensure maximum exhaust temperatures are kept at 50°C or lower at the back of the server.

Fan Speed Offset setting menu

Many of Dell's current-generation server products offer the ability to introduce a fan speed offset that causes fan speeds to increase over those calculated by the Dell Thermal Control algorithm. The Fan Speed setting menu allows the user to select Low Fan Speed Offset and High Fan Speed Offset. The Low Fan Speed Offset drives fan speeds to a moderate fan speed, and the High Fan Speed Offset drives fan speeds close to full speed. This setting may be useful when increased thermal margin is desired for custom high-power PCIe cards or for reducing system exhaust temperatures for adjacent equipment (like PDUs or switches).

Summary

The thermal control design of the current-generation of Dell PowerEdge servers allows a high level of optimization with newly supported usage modes, including support for ambient temperatures as high as 45°C. The advanced thermal control has increased capability to measure component temperature and power in addition to the new capability to manage subsystem power. This has allowed further improvements in tuning the thermal control response to individual configurations and usage with the benefit of reduced system power consumption and lower system acoustics.

Significant progress in Dell design also allows you to configure the server system thermal response to deployment needs. This includes the ability to adjust thermal control settings, optimize against maximum-performance and minimum-power needs, and increase system fan speed response to provide additional cooling.

Learn more

Visit [Dell.com](https://www.dell.com) for a more comprehensive overview of the latest generation of Dell PowerEdge servers.