

Custom Cooling Fan Options for Dell EMC PowerEdge Servers

Updated for the 14th Generation of Dell EMC Servers and iDRAC9

Abstract

Dell PowerEdge server fan control customization options and their potential use cases.

October 2019

Revisions

Date	Description
October 2019	Rev 1.0 – Updated release for iDRAC9 based Dell EMC Servers

Acknowledgments

This paper was produced by the following:

Authors: Hasnain Shabbir, Carlos Henry, Danny Whittington, and Jon Brown

Support: Dell EMC Server Thermal Controls Team

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © October 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [10/4/2019] [Technical White Paper] [[Manager]]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	4
1 Introduction.....	5
2 System Thermal Profile settings.....	6
2.1 Fan offset.....	6
2.2 Minimum Fan Speed (MFS)	7
2.3 Combined Settings – Fan offset and MFS	8
2.4 Maximum exhaust temperature limit.....	9
3 Customizing PCIe card cooling	11
4 Accessing thermal settings.....	12
4.1 Custom Fan Speed Options using HII (BIOS-based iDRAC settings)	13
4.1.1 System Thermal Profile Settings	16
4.2 Custom Fan Speed Options using iDRAC WebGUI.....	16
4.3 Custom fan speed options using RACADM	19

Executive summary

Sometimes, it is desirable to customize server cooling. For example, installing a third-party PCIe card could require more cooling beyond that which Dell EMC controls have provisioned. In other cases, the exhaust temperature of the server may be too high for a switch that is mounted behind the server. For these reasons, Dell EMC provides the fan control customization options that are discussed in this paper.

1 Introduction

Dell PowerEdge servers include proprietary thermal controls that automatically manage server component temperatures. Dell PowerEdge controls use sensor monitoring and knowledge of the specific hardware that is installed on the system to minimize fan speeds, airflow, and power consumption without sacrificing component reliability.

This white paper describes fan control customization options and suggests some potential use cases. Although the controls allow for cooling customization, the server does not allow fan speeds to drop below the threshold that is required to cool the server and any adapters that Dell provides. These options apply to server system fans and do not influence fans that are located in peripheral devices such as power supplies or PCIe cards.

NOTE: An iDRAC version of 3.30.30.30 or newer is recommended.

Some important points to keep in mind:

- All custom fan speed options are persistent. Once they are set and applied, they do not automatically change to the default setting during system reboot, power cycling, iDRAC, or BIOS updates.
- Some Dell EMC server platforms may not support some or any of these custom cooling options. If the options are not supported, they are not displayed, or you are not provided an opportunity to enter a custom value.
- Some features may be tied to an iDRAC license upgrade.

2 System Thermal Profile settings

System Thermal Profile settings can be changed based on the need to maximize performance or power efficiency.

- Default Thermal Profile Settings:
 - Implies that the thermal algorithm uses the same system profile setting that is defined under **System BIOS Settings→System Profile Settings** in the system BIOS menu.
 - Selecting any other option (Maximum Performance or Minimum Power) overrides the thermal settings that are associated with the system profile setting under **System BIOS Settings→System Profile Settings**.
- Maximum Performance:
 - Reduces probability of memory or CPU throttling
 - Increases probability of turbo mode activation
 - Generally, higher fan speeds at idle and stress loads
- Minimum Power:
 - Optimized for lowest system power consumption based on optimum fan power state
 - Generally, lower fan speeds at idle and stress loads
- Sound Cap:
 - Allows temporary deployment of servers in acoustic sensitive environment. System sets cap to CPU power to limit fan speed, consequently limiting acoustics. System performance is expected to be reduced as a result.

2.1 Fan offset

Selecting this option allows additional cooling of the server. A fan speed offset causes fan speeds to increase by the offset percentage value over baseline fan speeds that the Thermal Control algorithm calculates.

Possible values are:

- Low Fan Speed (+25%)—Drives fan speeds to a moderate fan speed.
- Medium Fan Speed (+50%)—Drives fan speeds close to medium.
- High Fan Speed (+75%)—Drives fan speeds close to full speed.
- Max Fan Speed (+100%)—Drives fan speeds to full speed.
- Off—Fan speed is set to Off, the default value. When set to Off, the percentage does not display. The default fan speed, according to internal thermal control algorithm, is applied with no additional offset. A fan speed offset set to Off should not be construed as the fan is not running or that fan speeds cannot change.

The fan speed offset option enables you to increase the system fan speed in four incremental steps. These steps are equally divided between the typical baseline speed and the maximum speed of the server system fans. Some hardware configurations result in higher baseline fan speeds, which then result in offsets other than the maximum offset to achieve maximum speed.

This option is used to provide additional cooling to custom configurations to the customer. The thermal algorithm maintains enough cooling without application of such offsets. The common use case is to add additional cooling to custom PCIe adapter cards, or to reduce the system exhaust temperature for platforms that do not support these specific customizations.

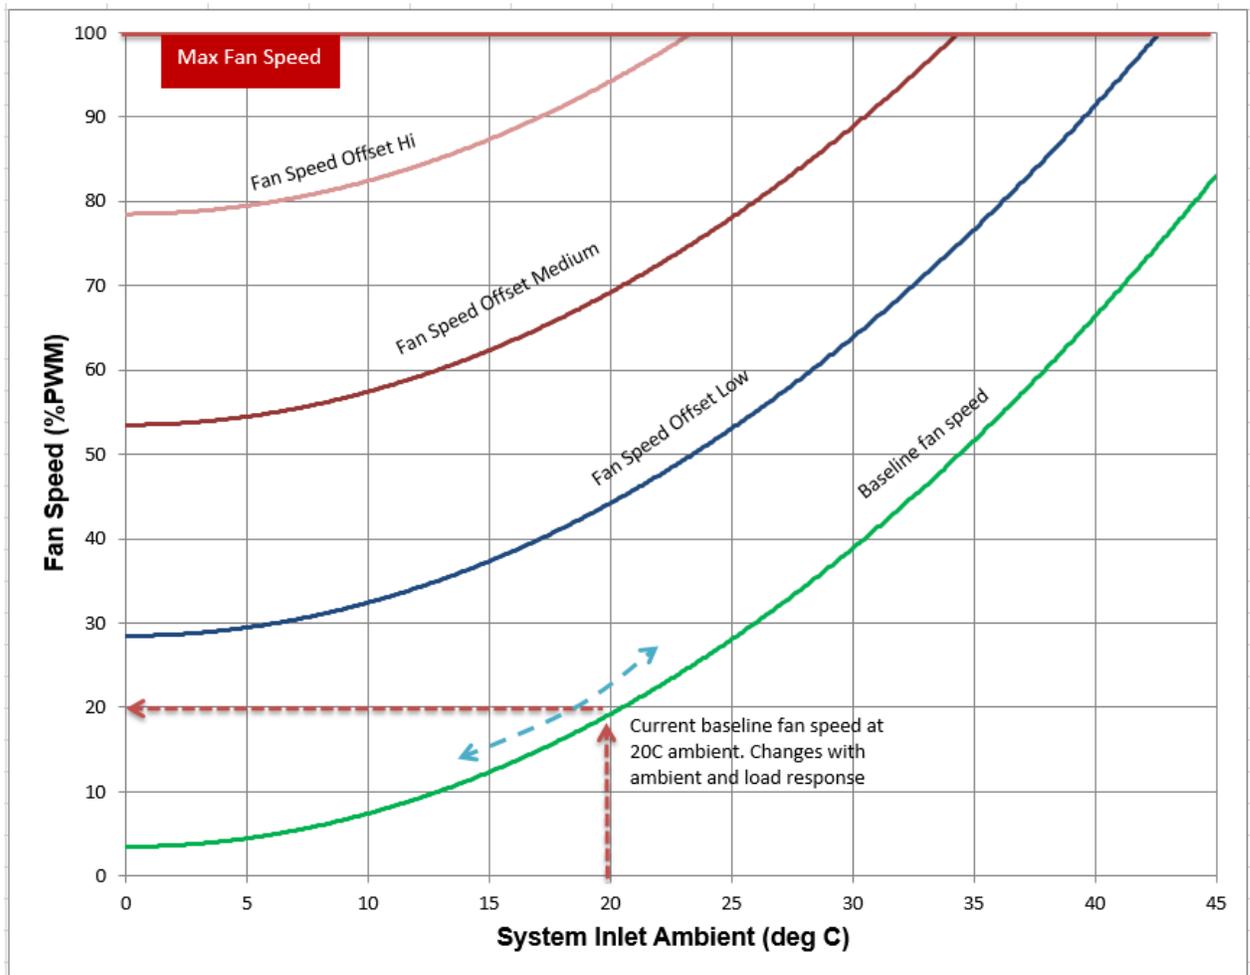


Figure 1 Illustration of how the fan speed offset increases the fan speed above the baseline (green line)

2.2 Minimum Fan Speed (MFS)

The MFS option allows the user to stipulate a lowest setting below which fans cannot drop. This setting is defined granularly as a percentage of the fan speed range, or % PWM.

System fans can run higher than the fan speed that the MFS option sets unless set to 100-percent, but not lower. For example, setting the MFS at 35-percent prevents the fan speed from dropping below 35% PWM.

NOTE: 0% PWM does not indicate that the fan is off. 0% PWM is the lowest fan speed that the fan can achieve.

The following figure illustrates how the MFS setting works along with the automatic controls:

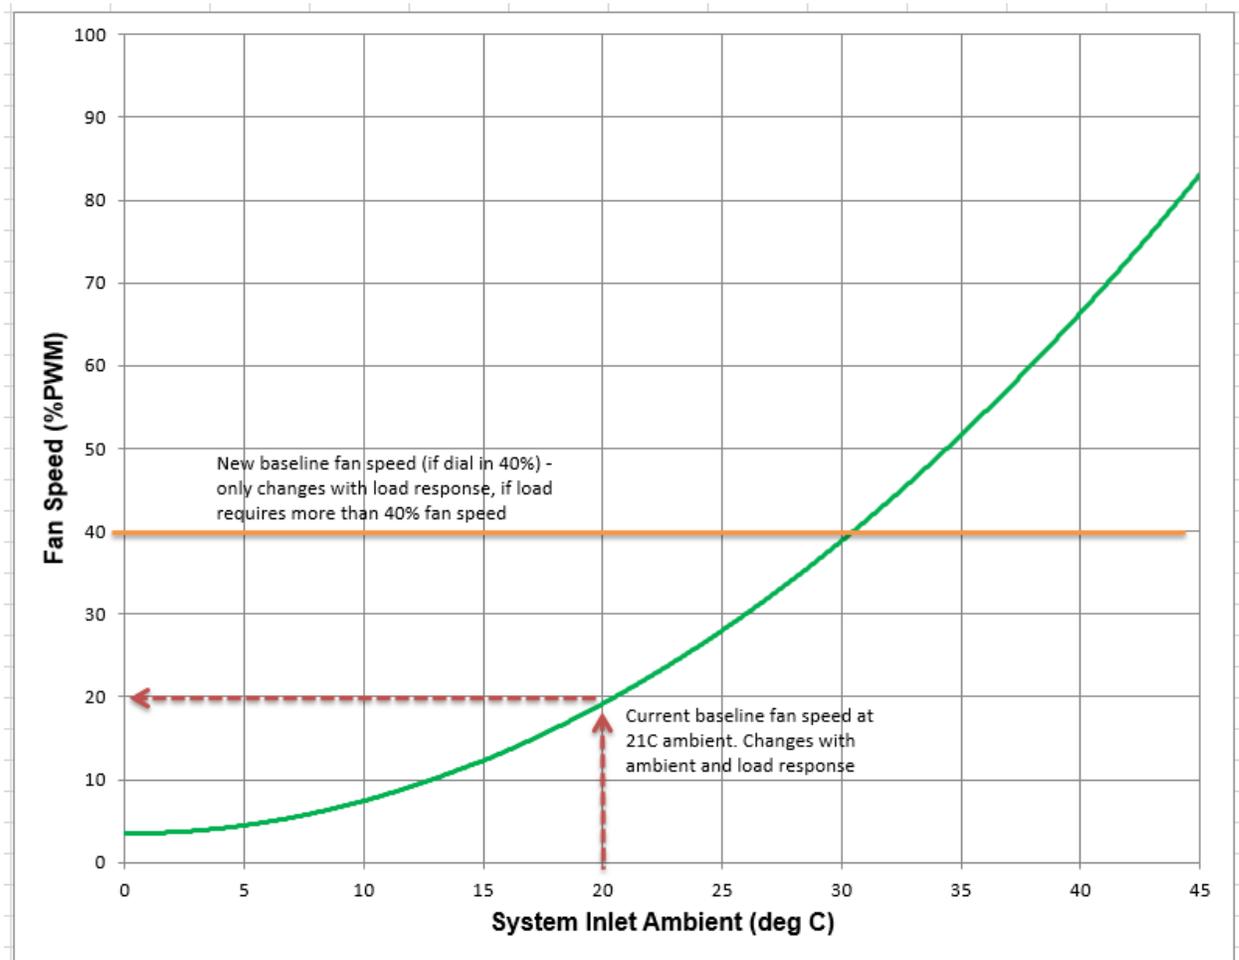


Figure 2 How the Minimum Fan Speed setting creates a fan speed floor, or lower bound

Other subsystems within the server, such as the CPU, can require higher fan speeds than the MFS or automatic baseline. Such requirements can occur during moderate to high use of that computing subsystem.

2.3 Combined Settings – Fan offset and MFS

Fan offset and MFS can be used together to achieve a combined effect. The following figure highlights how the two features might be used together to achieve a wanted result:

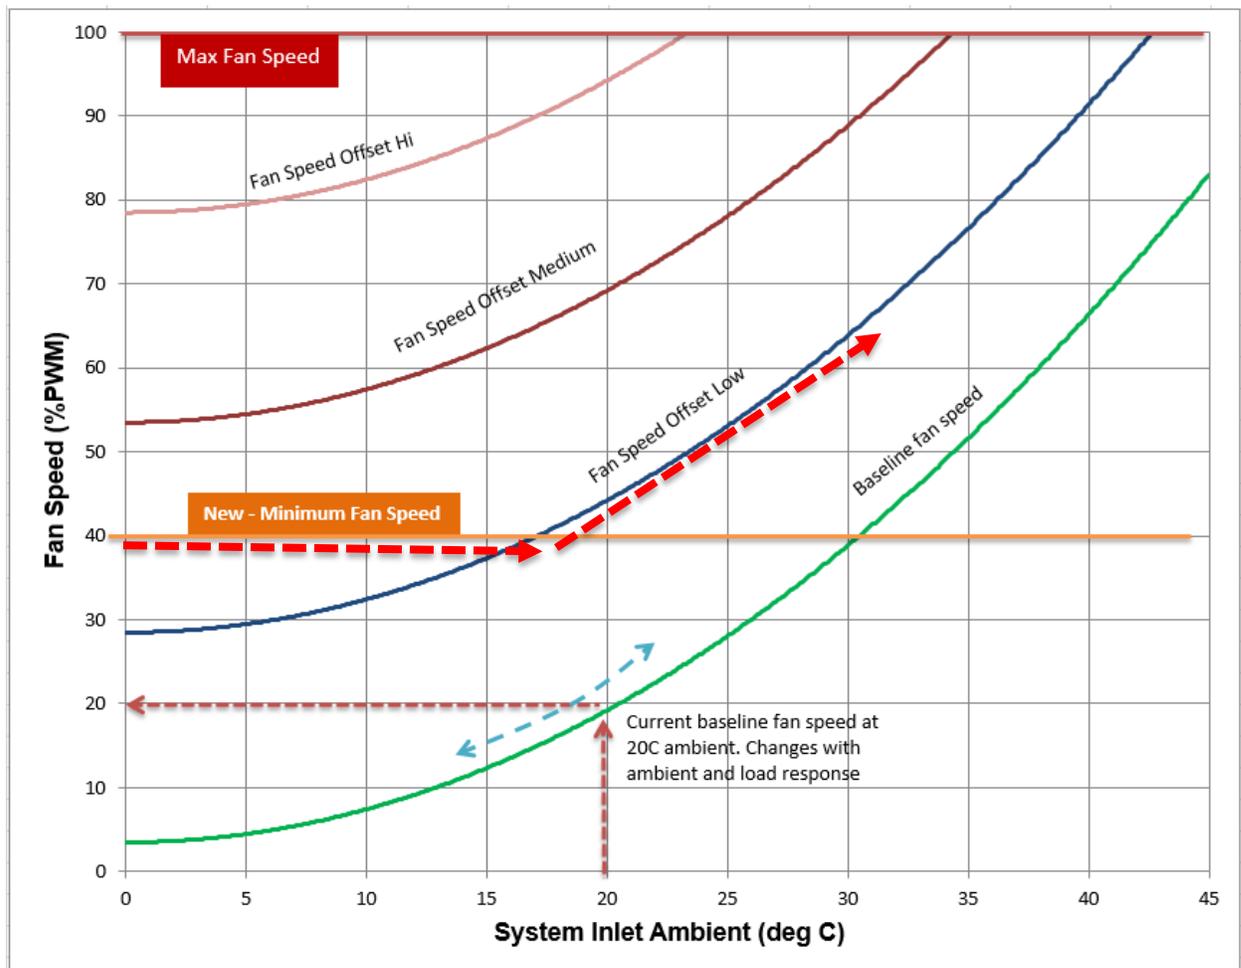


Figure 3 How options can be combined with MFS and Offset set simultaneously

If Low Offset and 40% MFS are applied, Fan Speed follows the red dotted line based on Inlet Ambient.

2.4 Maximum exhaust temperature limit

This option allows the system fan speeds to be regulated to prevent the system exhaust air temperature from exceeding the chosen threshold. Two important things to note regarding this option:

- Maintaining the wanted exhaust temperature cannot be guaranteed under all conditions. For example, if a system with a rich configuration of components is operating under full utilization and in a high ambient temperature environment, the system might not have enough cooling headroom to maintain a low exhaust temperature. The default value is 70°C (158°F) and is maintained under all conditions.
- It is also possible that there may not be any changes to fan speed because of this setting if the system configuration, inlet ambient conditions, and server workload do not result in higher exhaust ambient temperatures than the wanted setting.

Exhaust Temperature Setting	Resultant Relative Airflow
 70C (Default)	
 60C	
 55C	
 50C	

Figure 4 How airflow increases to maintain an exhaust temperature limit. Options go all the way to 40°C.

3 Customizing PCIe card cooling

In 14th Generation servers, Dell has implemented an industry-leading solution for PCIe cooling that enables the user better to understand the nature and amount of cooling to their PCIe cards. Since this topic is complex, see the white paper titled “PCIe Card Cooling with Dell PowerEdge Servers - Updates based on 14th Generation of Servers (iDRAC9)” dedicated to this subject.

4 Accessing thermal settings

Custom fan speed options described earlier can be applied on many different customer-facing interfaces such as iDRAC GUI, BIOS setup (F2), and RACADM. Figure 5 illustrates user interaction with server using various interfaces.

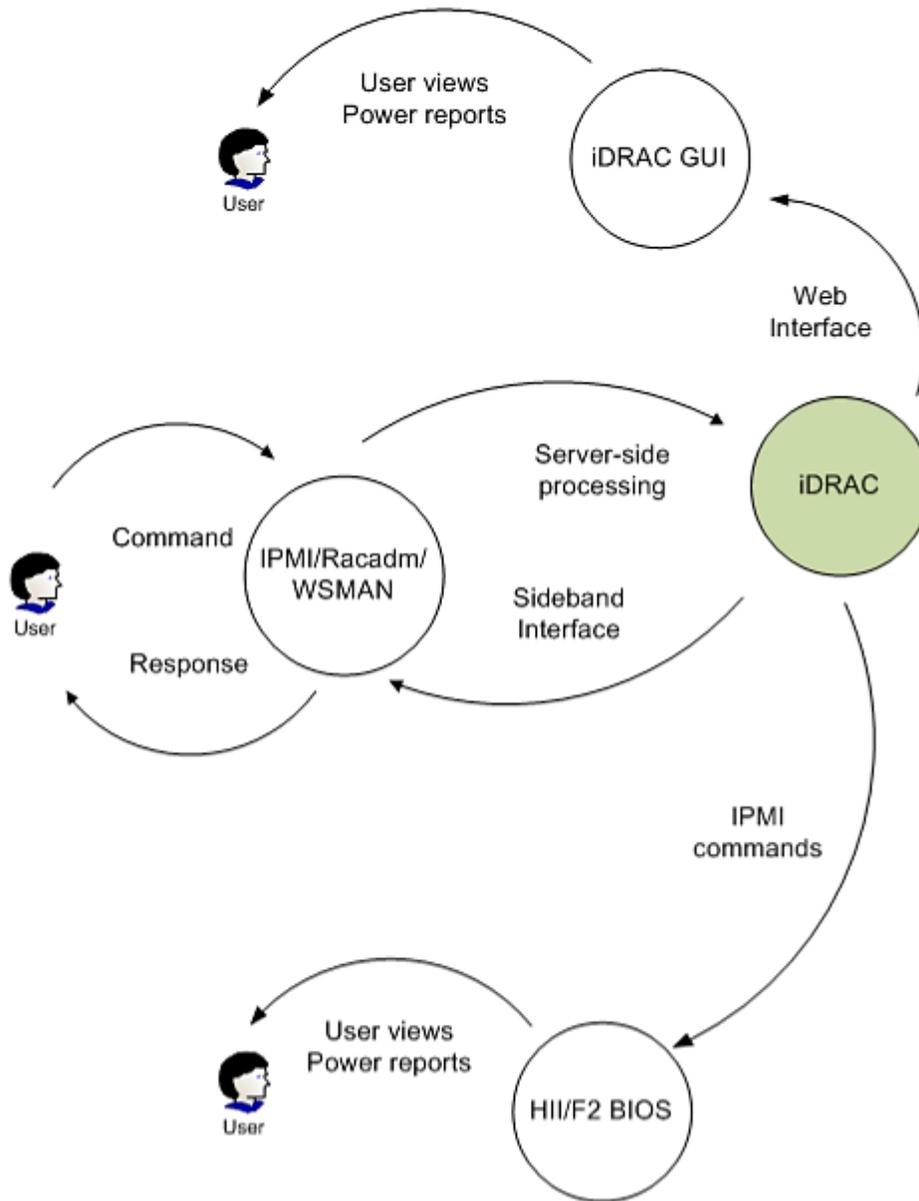


Figure 5 User paths to change thermal settings

The following is a summary of customer interfaces and typical reasons for using the interfaces for customization, in this case, fan speed customization:

- iDRAC UI
 - Web-browser based
 - Requires a network connection
 - UI-based and easy to navigate
 - No reboot required

NOTE: Alternate access using a USB A-A cable is also possible.

- Human Interface Infrastructure (HII) browser
 - Accessed through BIOS Setup (F2)
 - Does not require a network connection since setup is done directly on the server through BIOS setup
- Intelligent Platform Management Interface (IPMI)
- Web Service Management (WSMAN)
- Remote Access Controller Admin (RACADM)—Requires network connection and remote access to the server

4.1 Custom Fan Speed Options using HII (BIOS-based iDRAC settings)

The following screenshots describe the path to the BIOS setup screen—pressing F2 at boot and then selecting **iDRAC Settings**—to get to the custom fan speed options:

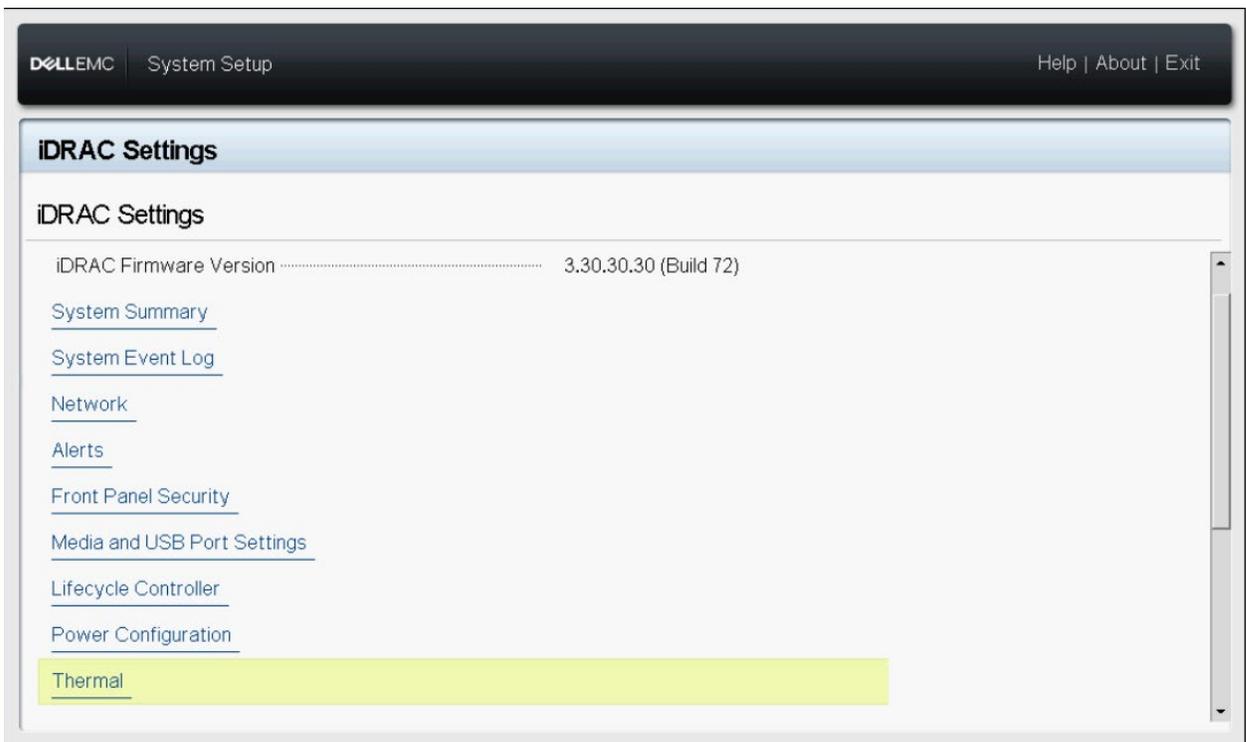


Figure 6 Thermal settings option within BIOS iDRAC settings

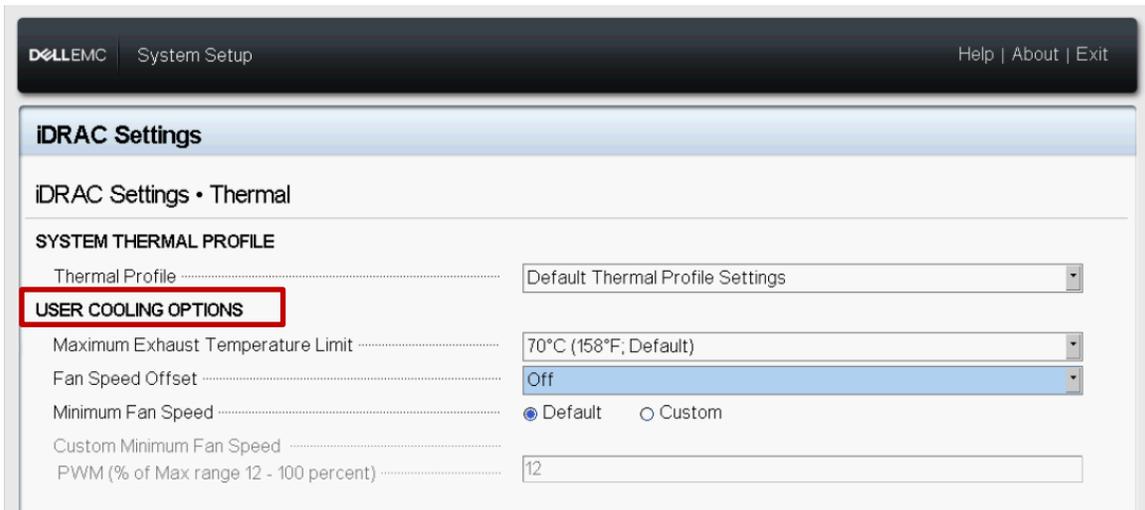


Figure 7 User cooling options section

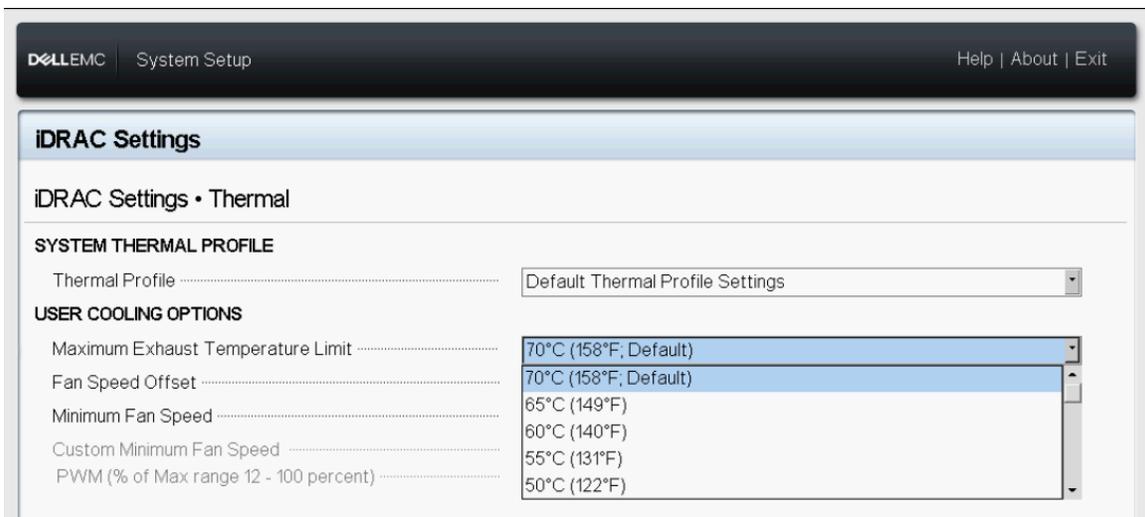


Figure 8 Setting Max Exhaust Temperature Limit

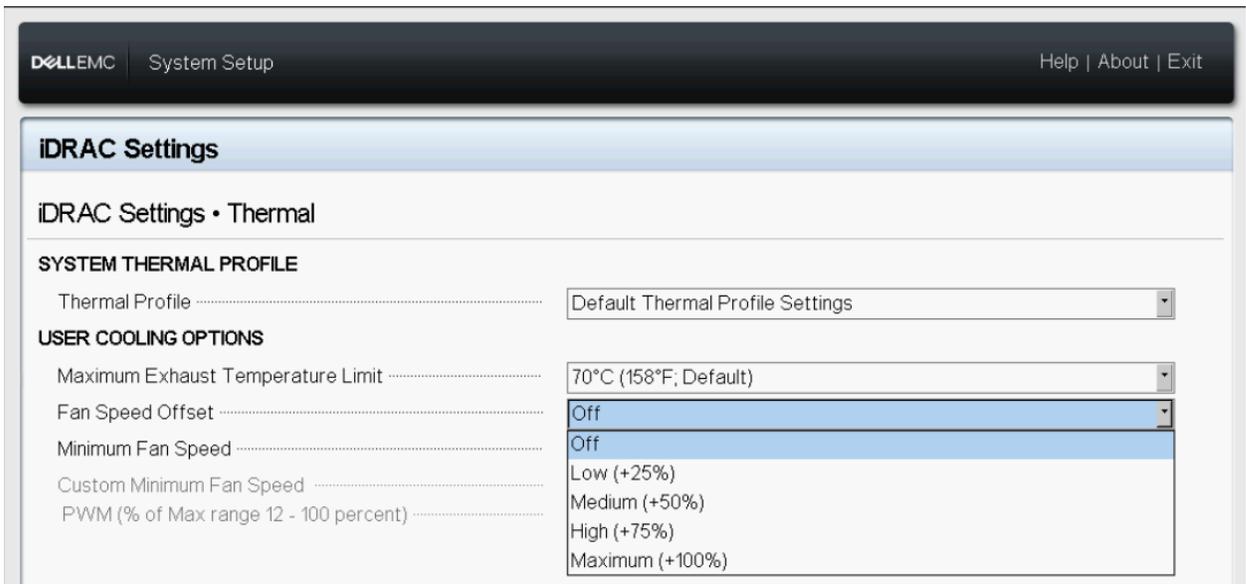


Figure 9 Setting Fan Speed Offset

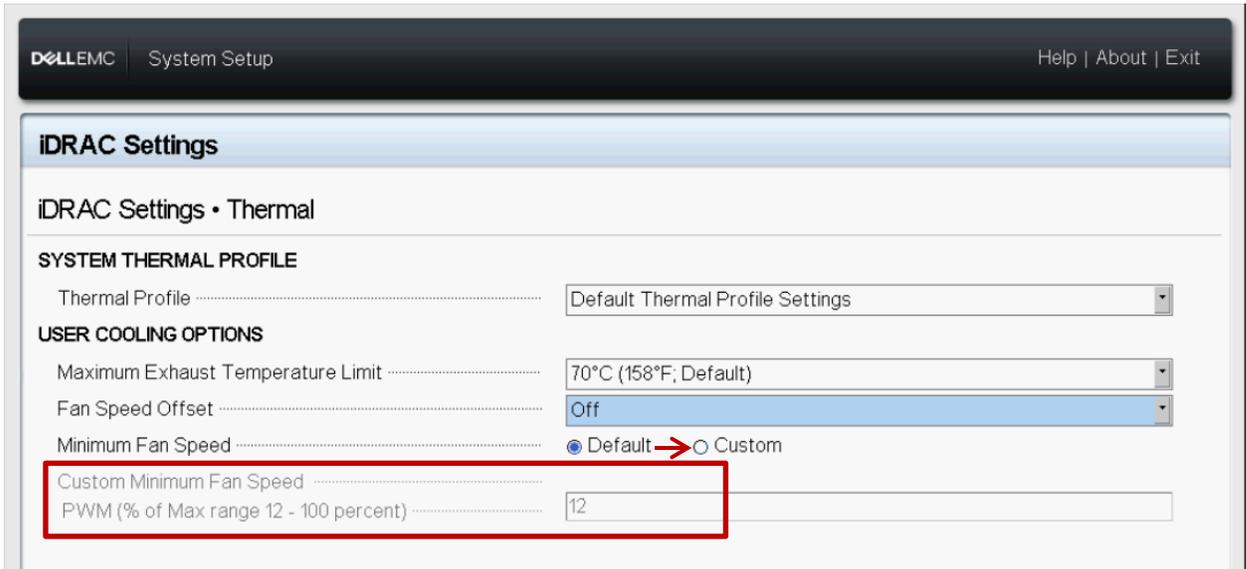


Figure 10 Setting Minimum Fan Speed (MFS)

4.1.1 System Thermal Profile Settings

In addition to the custom cooling options, the System Thermal Profile can be changed based on the desire to maximize performance or power efficiency. The Thermal Profile determines the settings that are used for the “Automatic Baseline” described in the previous sections.

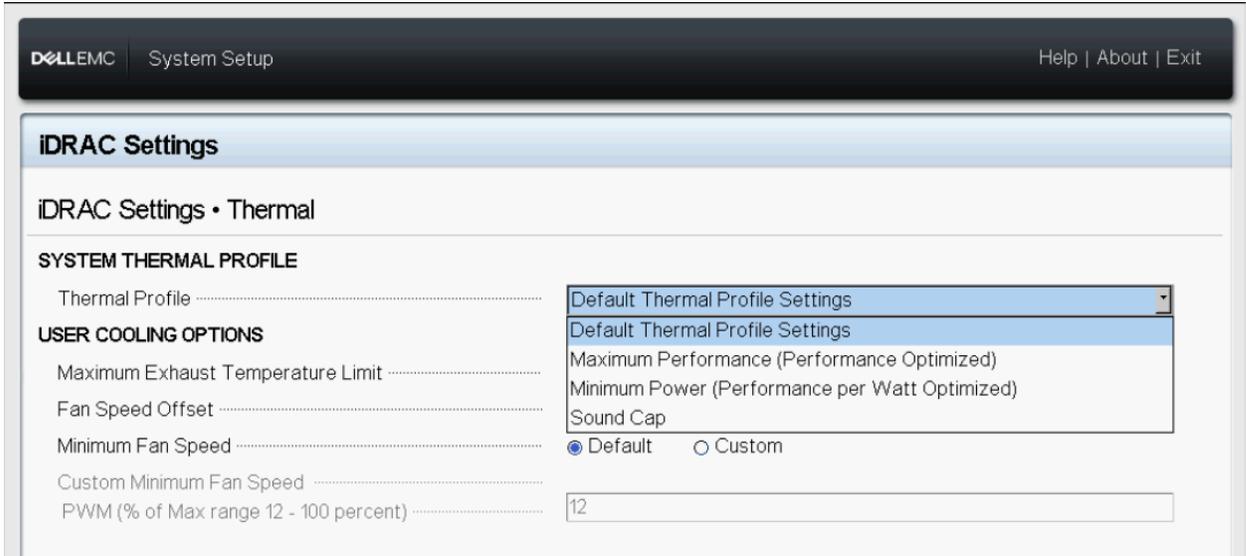


Figure 11 Setting thermal profile

4.2 Custom Fan Speed Options using iDRAC WebGUI

iDRAC WebGUI is a UI-based web management interface. You can apply custom fan settings using this interface as well. The screenshots that follow illustrate how to navigate to the options. You can connect to a server iDRAC WebGUI with an IP address on a remote web browser.

From the iDRAC Dashboard, select **System**→ **Overview**→ **Cooling** icon.

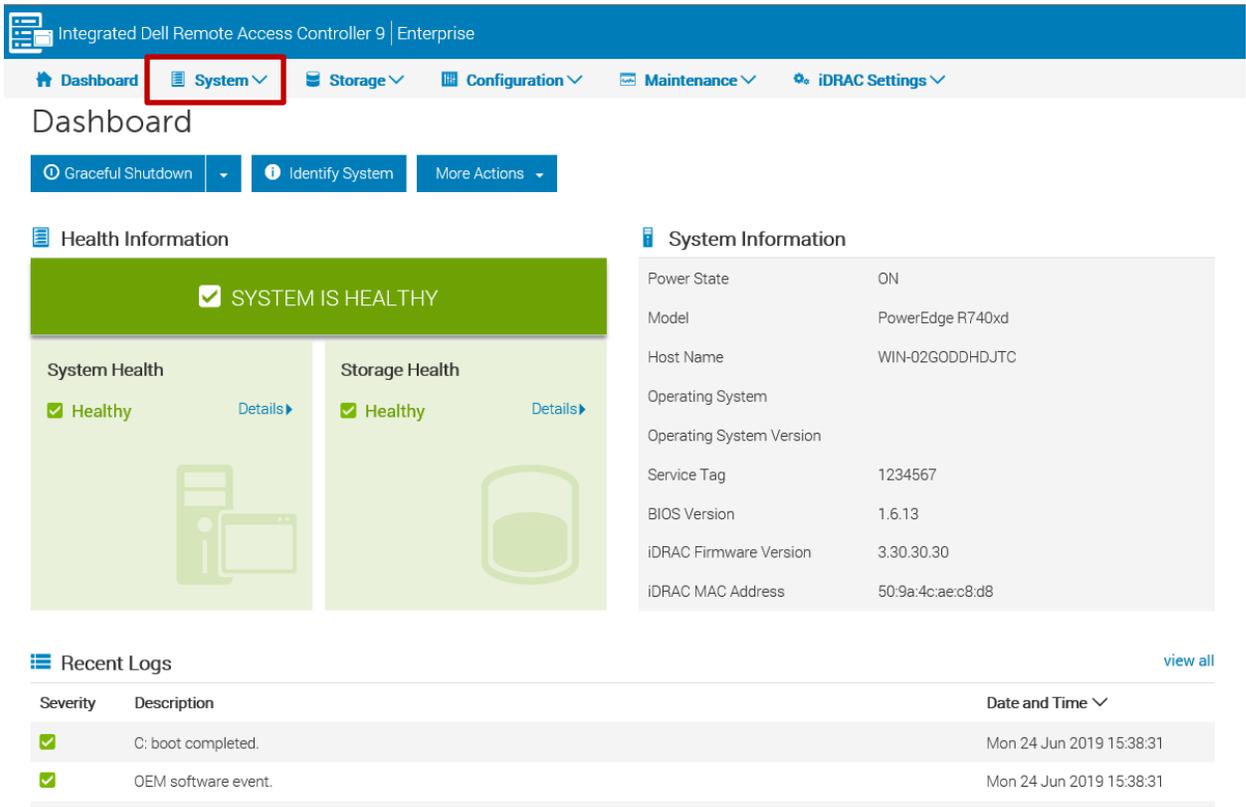


Figure 12 iDRAC dashboard view

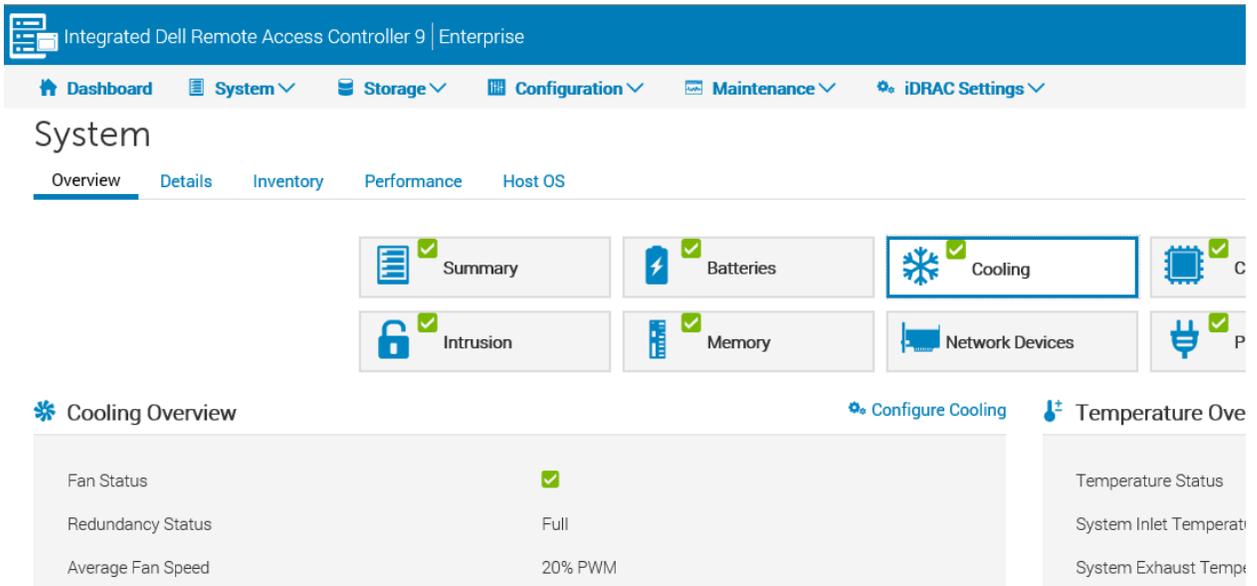


Figure 13 iDRAC cooling page

Then select **Cooling**→ **Configure Cooling**.

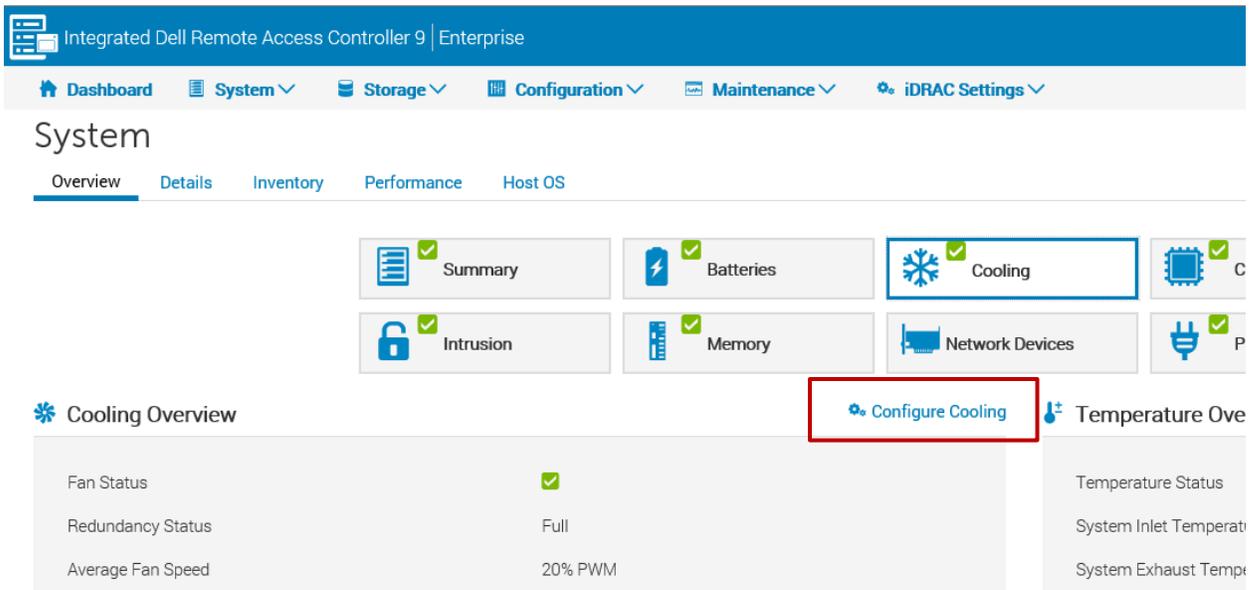


Figure 14 Configure cooling option

Make your selections under the **Cooling Configuration** option:

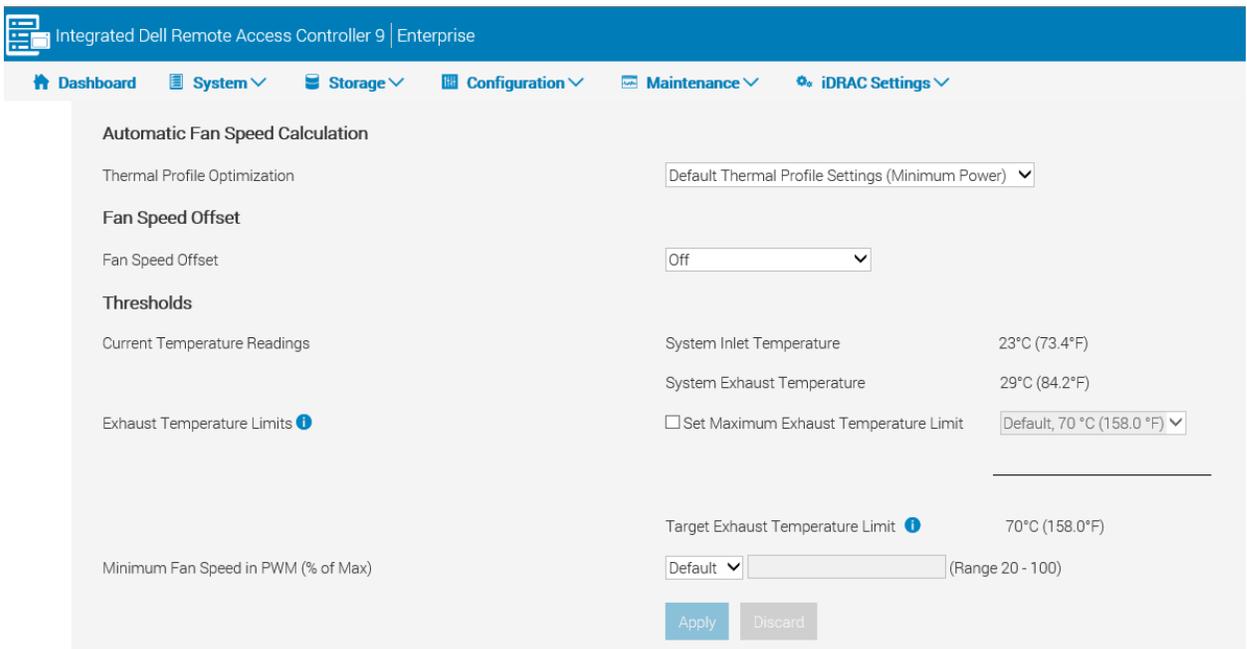


Figure 15 iDRAC custom cooling options

Each of the options that are described earlier can be toggled independently and set at the same time. For example, low Fan Speed Offset, and Exhaust of 50°C and Minimum Fan Speed of 20% can be set concurrently. The algorithm calculates the appropriate fan speed that meets all the customization requests. See the screenshots that follow for an example of such settings.

Hardware Settings

Cooling Configuration

Automatic Fan Speed Calculation

Thermal Profile Optimization: Maximum Performance (Performance Optimized)

Fan Speed Offset

Fan Speed Offset: Low Fan Speed (+25%)

Thresholds

Current Temperature Readings

System Inlet Temperature	23°C (73.4°F)
System Exhaust Temperature	29°C (84.2°F)

Exhaust Temperature Limits

Set Maximum Exhaust Temperature Limit: 60 °C (140.0 °F)

Target Exhaust Temperature Limit: 70°C (158.0°F)

Minimum Fan Speed in PWM (% of Max): Custom 35 (Range 20 - 100)

Apply Discard

Figure 16 Setting multiple custom control options

4.3 Custom fan speed options using RACADM

After logging into iDRAC, run `racadm get system.thermalsettings` to display the current settings. The results may vary depending on the iDRAC version used. It is recommended that you are working with at least 3.30.30.30 or newer release.

Here is an example of the command and its results:

```
get system.thermalsettings

[Key=system.Embedded.1#ThermalSettings.1]
#AirExhaustTempSupport=Supported
AirTemperatureRiseLimit=NO LIMIT
#CurrentSystemProfileValue=Minimum Power
#FanSpeedHighOffsetVal=75
#FanSpeedLowOffsetVal=25
#FanSpeedMaxOffsetVal=100
#FanSpeedMediumOffsetVal=50
FanSpeedOffset=Off
#MFSMaximumLimit=100
#MFSMinimumLimit=12
MinimumFanSpeed=255
SetMaximumExhaustTemperatureLimit=Disabled
#SystemCFMSupport=Supported
#SystemExhaustTemperature=35
#SystemInletTemperature=26
#SystemInletTemperatureSupportLimitPerConfiguration=35
#TargetExhaustTemperatureLimit=70
ThermalProfile=Default Thermal Profile Settings
```

The following table can be used as a guide to configure these settings using the `racadm system.thermalsettings` command:

Table 1 Configure settings using `racadm system.thermalsettings` command

Object	Description	Usage	Example
AirExhaustTemp	<ul style="list-style-type: none"> • Max Air Exhaust Temperature limit setting. • Set Exhaust Temperature Limit to any one value (platform dependent), where 0,1,2,3,4 and 255 correspond to 40,45,50,55,60°C and 70°C (default). • 	Values from 0-4 and 255 (default)	<p>You can run <code>racadm help</code> on each attribute to show the available options. For example:</p> <pre>/home/root# racadm help system.thermalsettings.airexhausttemp</pre> <p>AirExhaustTemp—Specifies the Average Air Exhaust Temperature Usage: 0- 40; 1- 45; 2- 50; 3- 55; 4- 60; 5- 65; 255- 70; Default – 255 Dependency -- Feature must be supported to set the value.</p> <p>To check existing setting on the system: <code>racadm get system.thermalsettings.AirExhaustTemp</code> It reports: AirExhaustTemp=70 This means that system is set to limit air exhaust temperature to 70°C.</p> <p>To set exhaust temperature limit to 60°C: <code>racadm set system.thermalsettings.AirExhaustTemp 4</code> The result is: "Object value modified successfully," and a 60°C exhaust limit is set.</p> <p>If a certain AirExhaust Temperature limit (such as 40°C) is not supported by the platform, then running the command: <code>racadm set system.thermalsettings.AirExhaustTemp 0</code> results in error: ERROR: RAC947: Invalid object value specified.</p> <p>Ensure that you specify the value depending on the type of object. For more information, see RACADM help. To set the limit to default value: <code>racadm set system.thermalsettings.AirExhaustTemp 255</code></p>

<p>FanSpeedHighOffsetVal</p>	<ul style="list-style-type: none"> • Getting this variable reads the fan speed offset value in %PWM for High Fan Speed Offset setting. • This value is platform-dependent. • Use FanSpeedOffset to set this value using index value 1. 	<p>Values from 0-100</p>	<p>racadm get system.thermalsettings FanSpeedHighOffsetVal</p> <p>This returns a value like “75” meaning that when you use the command “racadm set system.thermalsettings FanSpeedOffset 1”, it applies a fan speed offset of High (75% PWM) over baseline fan speed.</p>
<p>FanSpeedLowOffsetVal</p>	<ul style="list-style-type: none"> • Getting this variable reads the fan speed offset value in %PWM for Low Fan Speed Offset setting. • This value is platform-dependent. • Use FanSpeedOffset to set this value using index value 0. 	<p>Values from 0-100</p>	<p>racadm get system.thermalsettings FanSpeedLowOffsetVal</p> <p>This returns a value like “25,” meaning that when you use command “racadm set system.thermalsettings FanSpeedOffset 0”, it applies a fan speed offset of Low (25% PWM) over baseline fan speed.</p>
<p>FanSpeedMaxOffsetVal</p>	<ul style="list-style-type: none"> • Getting this variable reads the fan speed offset value in %PWM for Max Fan Speed Offset setting. • This value is platform-dependent. • Use FanSpeedOffset to set this value using index value 3. 	<p>Values from 0-100</p>	<p>racadm get system.thermalsettings FanSpeedMaxOffsetVal</p> <p>This returns a value like “100,” meaning that when you use command “racadm set system.thermalsettings FanSpeedOffset 3”, fan speed offset of Max (meaning full speed, 100% PWM) is applied. Usually, this offset results in fan speeds going to full speed.</p>
<p>FanSpeedMediumOffsetVal</p>	<ul style="list-style-type: none"> • Getting this variable reads the fan speed offset value in %PWM for Medium Fan Speed Offset setting. • This value is platform-dependent. • Use FanSpeedOffset to set this value using index value 2. 	<p>Values from 0-100</p>	<p>racadm get system.thermalsettings FanSpeedMediumOffsetVal</p> <p>Returns a value (for example) “50,” meaning that when you use command “racadm set system.thermalsettings FanSpeedOffset 2”, a fan speed offset of Medium (50% PWM) over baseline fan speed is applied.</p>

<p>FanSpeedOffsetI</p>	<ul style="list-style-type: none"> Using this with “get” reports the existing Fan Speed Offset setting. Using this with “set” allows setting the wanted fan speed offset value. The index governs what offset is applied and FanSpeed...Val variables (defined earlier) are the values at which the offsets are applied. 	<p>0-Low Fan Speed, 1-High Fan Speed, 2-Medium Fan Speed, 3-Max Fan Speed, 255-None</p>	<p>Example with get:</p> <pre>racadm get system.thermalsettings.FanSpeedOffset</pre> <p>Reports current setting. Example with set:</p> <pre>racadm set system.thermalsettings.FanSpeedOffset 1</pre> <p>Results in fan speed offset set to High value (as defined in “FanSpeedHighOffsetVal”).</p>
<p>MFSMaximumLimit</p>	<ul style="list-style-type: none"> Read Maximum limit for MFS.using index value 1. 	<p>Values from 1 - 100; Default – 255 (meaning None)</p>	<p>Example:</p> <pre>racadm get system.thermalsettings.MFSMaximumLimit</pre> <p>Reports the highest value that can be set using MinimumFanSpeed option.</p>
<p>MFSMinimumLimit</p>	<ul style="list-style-type: none"> Read Minimum limit for MFS. 	<p>Values from 0 - MFSMaximumLimit; Default – 255 (meaning None)</p>	<p>Example:</p> <pre>racadm get system.thermalsettings.MFSMinimumLimit</pre> <p>Reports the lowest value that can be set using MinimumFanSpeed option.</p>
<p>MinimumFanSpeed</p>	<ul style="list-style-type: none"> Allows configuring of Minimum Fan speed that is wanted for the system to operate. Defines the baseline (floor) value for fan speed, and system enables fans to go below this defined fan speed value. This value is %PWM value for fan speed. 	<p>Values from MFSMinimumLimit to MFSMaximumLimit When get command reports 255, it means no user configured offset applied.</p>	<p>Example:</p> <pre>racadm set system.thermalsettings.MinimumFanSpeed 45</pre> <p>Forces the system minimum speed to not go below 45% PWM (45 should be a value between what is read from MFSMinimumLimit to MFSMaximumLimit).</p>

<p>ThermalProfile</p>	<ul style="list-style-type: none"> • Selection for Thermal Profile • Allows system profile to be set at preferred setting for thermal behavior that is associated with the profile. 	<p>0-Default Thermal Profile, 1-Maximum performance, 2-Minimum Power, 3-Sound Cap (on supported platforms).</p>	<p>Example using get:</p> <pre>racadm get system.thermalsettings.ThermalProfile</pre> <p>Reports the current ThermalProfile setting</p> <p>Example using set:</p> <pre>racadm set system.thermalsettings.ThermalProfile 1</pre> <p>Results in Thermal Profile set to Maximum Performance Mode.</p>
-----------------------	---	---	--