

Strategically Manage Power in PowerEdge Server PSUs Using Power Settings in iDRAC

Abstract

This technical whitepaper describes the server power settings and their effects on PowerEdge server power management and control. Dell incorporates power margin for system inventory, telemetry of subsystem components, and control features into a cohesive solution. This whitepaper also includes instructions to configure the settings using Redfish URI.

August 2025

Revisions

Date	Description
August 2025	Initial release

Acknowledgments

Author: Craig Klein, Vaishnavi S

Support: Samiksha Aggarwal, Technical Content Developer

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.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © 2025 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [8/7/2025] [Technical Whitepaper] [670]

Contents

Revisions.....	2
Acknowledgments.....	2
Contents.....	3
1 Power Supply Configuration.....	4
2 iDRAC settings for Fault Tolerance.....	5
3 Leverage iDRAC redfish to Get or Set.....	6
4 Helpful Links.....	7

1 Power Supply Configuration

System power management enables Power Supply Unit (PSU) capacity optimization, enhancing the effectiveness of system components. The PSU's power capacity is selected to support sustained workloads at full performance. However, PSUs with excessive power capacity may operate at lower efficiency, leading to energy waste and reduced operational cost-effectiveness.

From a power perspective, a redundant system has enough total PSU capacity such that some non-zero number of operational PSUs can become non-operational and the system will continue to operate. You select the target redundancy, which determines when and how a loss of redundancy is reported to you. The system's real-time calculated redundancy is based on the calculated power budget and the total power capacity of the available PSUs.

Table 1 Dell Power Supply Redundancy Types

Type	Description	Performance
Non-Redundant (N+0)	The system's maximum expected power consumption remains within the total power available from the PSU(s), provided all configured PSUs are healthy and active.	If a PSU is lost, the system will continue to operate if the remaining PSUs can handle the power requirement. It is not guaranteed that the system will maintain the same operational conditions. If the load cannot be handled, the server will be powered off.
PSU Redundant (N+M)	The system's peak power consumption does not exceed the power available from the PSU(s), even in scenarios where up to M-1 PSUs are lost.	The system is guaranteed to maintain the operational condition until M PSU(s) stop functioning. The system may throttle to stay online even in such extreme conditions.
A/B or Grid Redundant (N+N)	For maximum redundancy, customers will have multiple grids. Grid A will supply power to half of the PSUs and Grid B will supply to the other half.	If Grid A loses power, PSUs in Grid B will maintain the operational conditions.

Note: The system's maximum power consumption—including short-term excursions—does not exceed the power available from the PSU(s), even if up to M-1 PSUs become non-operational.

iDRAC performs power budgeting when booting the system which determines the minimum redundancy configuration required for maximum performance. The iDRAC interface allows the customer to configure the type of redundancy required. If the redundancy calculated by iDRAC is less than the redundancy configured by you, then a message is logged in Lifecycle Controller indicating that the redundancy state is lost.

2 iDRAC settings for Fault Tolerance

The CPU subsystem can be fully throttled by using the `PROCHOT#` signal, and the PCIe subsystem can be fully throttled using the `PWRBRK#` signal. These signals drive their respective subsystems to the minimum supported frequency and impacting performance. If a PSU stops functioning, then the fault-tolerance feature can be adjusted to mitigate heavy throttling using Fault Tolerant Redundancy through iDRAC settings.

iDRAC power management uses a combination of two types of power control and limitation:

- **Preemptive peak limiting**—Prevents power transients from reaching the power source by limiting the maximum instantaneous CPU power consumption. This value is lesser than the threshold value based on the initial power budgeting values and the PSU Extended Power range (EPR).
- **Runtime limiting**—Uses a reactive control loop to limit the average CPU power consumption. Because the control loop relies on power readings to adjust the CPU power limits, a short lag can exist between any transients and the resultant compensation in the CPU. Because the control loop can be adjusted on the fly when any extra power is available, it prevents stranded power and maximizes the overall system performance when it is power constrained. This method may allow some transients to reach the power source.

A fully redundant configuration continues to operate at the same level of performance after the loss of redundancy without applying power limits. However, further loss of PSUs may require power limiting to counter any shortage of available power in the reactive control loop.

You can disable the `PROCHOT#` and `PWRBRK#` features using non-volatile iDRAC attributes to avoid potential negative impacts on streaming or high-performance computing applications. However, disabling using the non-volatile iDRAC attributes may result in automatically shutting down the server.

Table 2 Fault Tolerance based on the User Redundancy Selection and iDRAC Settings

Type	Description
Non-Redundant	Disables <code>PROCHOT#</code> and <code>PWRBRK#</code>
Redundant (N+M) or (N+N)	Enables <code>PROCHOT#</code> and <code>PWRBRK#</code>

3 Leverage iDRAC Redfish to Get or Set

After setting the following attributes, configuring the server PSU redundancy configuration to non-redundant disables `PROCHOT` and `PWRBRK` throttling:

1. **DisableUnderBudgetThrottling**—Manages throttling when there is insufficient available power.
2. **InputLossThrottling**—Manages throttling when there is a loss of input power to the PSU.

Redfish URI :

```
/redfish/v1/Managers/iDRAC.Embedded.1/Oem/Dell/DellAttributes/System.Embedded.1
```

Run the GET operation on this URI for the attributes.

```
{
  "Attributes": {
    "ServerPwr.1.DisableUnderBudgetThrottling": "0",
    "ServerPwr.1.InputLossThrottling": "1"
  }
}
```

PATCH on the same uri with this payload...

```
{
  "Attributes": {
    "ServerPwr.1.DisableUnderBudgetThrottling": "value",
    "ServerPwr.1.InputLossThrottling": "value"
  }
}
```

4 Helpful Links

- Dell iDRAC Attribute Registry:
 - [Integrated Dell Remote Access Controller 9 Attribute Registry | Dell US](#)
- Full redundancy and fault-tolerant redundancy
 - <https://infohub.delltechnologies.com/it-it/p/full-redundancy-vs-fault-tolerant-redundancy-for-poweredge-server-psus/>