



# Power Efficiency “How To” for the Dell PowerEdge Server Portfolio

How to order and configure Dell PowerEdge servers for optimal power efficiency

Dell Enterprise Systems Engineering  
March 2016

John Jenne  
Craig Klein

## Revisions

Date	Description
March 2016	Initial release

This white paper 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. Copyright © 2016 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



## Table of contents

Executive summary .....	5
Introduction .....	5
How to order a server for power efficiency .....	5
Advisor tools .....	5
ESSA .....	5
Processor and Memory Selector .....	7
Order validator .....	22
Components.....	23
Processor .....	23
Memory .....	23
Storage .....	23
Power supply .....	24
BIOS Settings .....	25
How to configure a server for power efficiency .....	25
No performance impact.....	25
Hardware setup.....	25
Unused services and ports .....	25
Power supply.....	25
Component updates .....	26
Potential for minor performance impact.....	26
System profiles.....	26
Power management .....	27
Reduce top-end power consumption .....	29
Thermal management.....	29
Miscellaneous.....	30
Power capping .....	30
Operating systems.....	31
About the authors .....	31

## Tables

Table 1	Setup options: Unused devices and ports .....	25
Table 2	Setup options: Power supply.....	26
Table 3	Setup options: System profile .....	27
Table 4	System profile definitions.....	27
Table 5	Setup options: Power management .....	28
Table 6	Reducing top-end power consumption .....	29



Table 7	Setup options: Thermal management.....	30
Table 8	Miscellaneous power setup options .....	30
Table 9	Setup options: Power capping.....	30

## Figures

Figure 1	PSU Help Me Choose example recommending the 750W PSU .....	6
Figure 2	Dell PowerEdge Processor and Memory Selector - Welcome page .....	7
Figure 3	Dell PowerEdge Processor and Memory Selector - Configuration page .....	8
Figure 4	Dell PowerEdge Processor and Memory Selector - Server selection.....	9
Figure 5	Dell PowerEdge Processor and Memory Selector - Processor information .....	10
Figure 6	Dell PowerEdge Processor and Memory Selector - Processor details.....	11
Figure 7	Dell PowerEdge Processor and Memory Selector – Setting processor count .....	12
Figure 8	Dell PowerEdge Processor and Memory Selector – Setting target memory size.....	13
Figure 9	Dell PowerEdge Processor and Memory Selector – Setting processor performance bin...	14
Figure 10	Dell PowerEdge Processor and Memory Selector – Memory reliability features.....	15
Figure 11	Dell PowerEdge Processor and Memory Selector – Target memory size invalid.....	16
Figure 12	Dell PowerEdge Processor and Memory Selector – Target memory size correction .....	17
Figure 13	Dell PowerEdge Processor and Memory Selector – Submit data .....	18
Figure 14	Dell PowerEdge Processor and Memory Selector – Results page for unbalanced configuration .....	19
Figure 15	Dell PowerEdge Processor and Memory Selector – Results page for final configuration ..	20
Figure 16	Dell PowerEdge Processor and Memory Selector – Get report.....	21
Figure 17	Dell PowerEdge Processor and Memory Selector – Report template .....	22
Figure 18	Order validator — Power supply error message .....	22
Figure 19	PSU power dissipation vs. system load.....	24
Figure 20	Power capping example.....	31



## Executive summary

As energy costs rise and the concerns surrounding the environmental impact of energy generation also increase, the focus on power efficiency within data centers is more important than ever. IT departments are becoming increasingly concerned with power efficiency, looking to save every watt possible without sacrificing the system performance required to meet their business needs. At first glance, optimizing servers for power efficiency may seem like a daunting task due to the seemingly overwhelming complexity of power management features, the multitude of configuration options, and the potential side effects of enabling these features. To help customers navigate through their system power options, Dell has focused on simplifying power management on the Dell™ PowerEdge™ 13th generation server portfolio by providing easy-to-use system profiles for those customers that want fast and simple deployment. For those customers wishing to fine-tune their systems for the optimal performance/power balance appropriate for their particular workloads, Dell also incorporates flexible and highly granular configuration options with detailed help windows, giving Dell customers the tools to custom tailor their servers on a system-by-system basis.

## Introduction

The Dell PowerEdge 13<sup>th</sup> generation servers are designed with a focus on maximizing performance while minimizing power consumption. To minimize power consumption in the data center, you must select components that emphasize power efficiency, as well as configure settings for efficiency, such as BIOS and iDRAC, before deployment. This white paper serves as a “how to” guide for optimizing a PowerEdge server for power consumption. The majority of this guide applies to both rack and tower servers.

## How to order a server for power efficiency

Dell PowerEdge servers provide many configuration options that you can use to minimize a server’s power consumption. Dell offers tools to help you determine the server configurations that best fit your needs and reduce power use.

## Advisor tools

Dell offers online advisor tools to help you navigate through some of the complex component selections. These tools quickly and easily help you assess and determine which PowerEdge configuration options best suit your needs.

## ESSA

The Dell [Energy Smart Solution Advisor \(ESSA\)](#)<sup>1</sup> offers a feature called **Help Me Choose** to assist with the selection of a power supply for PowerEdge rack and tower servers. ESSA can be used to estimate the power requirements for a particular server configuration, to estimate available headroom for a desired power supply, and to assess the return-on-investment benefit for right-sizing. With the **Help Me Choose** feature for PSUs, you can choose from PSUs with different capacity points, efficiency ratings, input types, and redundancy configurations to find the best PSU solution to meet your needs.

Figure 1 illustrates how **Help Me Choose** for PSUs helps guide you when selecting a power supply for a particular configuration. The lowest supported PSU capacity is the default recommendation, but you can also

---

<sup>1</sup> <http://essa.us.dell.com/dellstaronline/Launch.aspx/ESSA?c=us&l=en&s=corp>



choose to upgrade to a larger capacity PSU for additional headroom or to a PSU with higher efficiency, such as the 1100W Platinum+ or the 750W Titanium power supply.

**Help Me Choose**
?
✕

PowerEdge R730
Cost Per kWh: \$0.14/kWh

PSU Option	Estimated Input Power At Workload (W)	Estimated Maximum Potential Power (W)	Estimated Annual Energy Cost(Continual Operation)	
			At Workload (\$USD)	At Maximum Potential Power (\$USD)
<input type="radio"/> Dual, Hot-plug, Redundant Power Supply (1+1), 495W	820	1035	1005.65	1269.32
<input checked="" type="radio"/> Dual, Hot-plug, Redundant Power Supply (1+1), 750W <span style="float: right; font-size: small;">Dell Recommended</span>	820	1035	1005.65	1269.32
<input type="radio"/> Dual, Hot-plug, Redundant Power Supply (1+1), 750W, Titanium, 200-240VAC	820	1035	1005.65	1269.32
<input type="radio"/> Dual, Hot-plug, Redundant Power Supply (1+1), 1100W	798	1018	978.67	1248.48
<input type="radio"/> Single, Hot-plug Power Supply (1+0), 495W	820	1035	1005.65	1269.32
<input type="radio"/> Single, Hot-plug Power Supply (1+0), 750W	820	1035	1005.65	1269.32

Warning: Exceeding PSU rated capacity may drive system throttling

Figure 1 PSU Help Me Choose example recommending the 750W PSU



## Processor and Memory Selector

The [Dell PowerEdge Processor and Memory Selector](http://www.poweredgecpumemory.com/)<sup>2</sup> is an online tool that provides advice for selecting a memory configuration that is optimized for price, performance, power, and/or RAS based on user input. Since the memory controller is integrated into an Intel® Xeon® processor, selecting the processor is also a crucial part of building an optimal memory configuration. Dell PowerEdge servers support a rich assortment of memory configurations, and with the PowerEdge Processor and Memory Selector, you can easily select a configuration to meet your workload requirements.

The first page in the Processor and Memory Selector (Figure 2) is the Welcome page that explains how to use the tool. To begin, select a processor type, and then click **Next** to go the Configuration page.

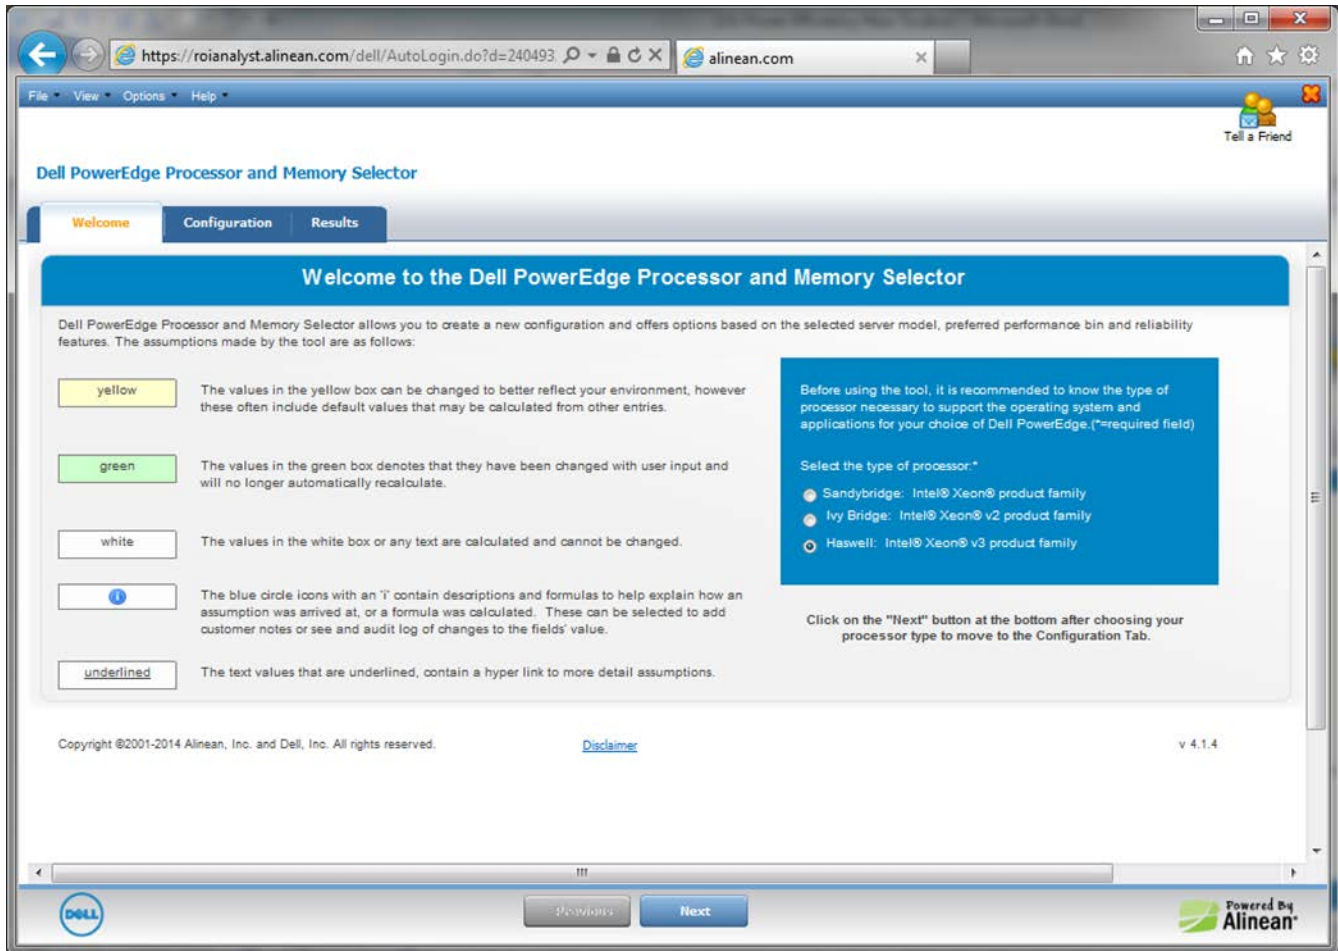


Figure 2 Dell PowerEdge Processor and Memory Selector - Welcome page

<sup>2</sup> <http://www.poweredgecpumemory.com/>



On the Configuration page (Figure 3), enter information the tool can use to provide configuration guidance. In Step 1, select a server type.

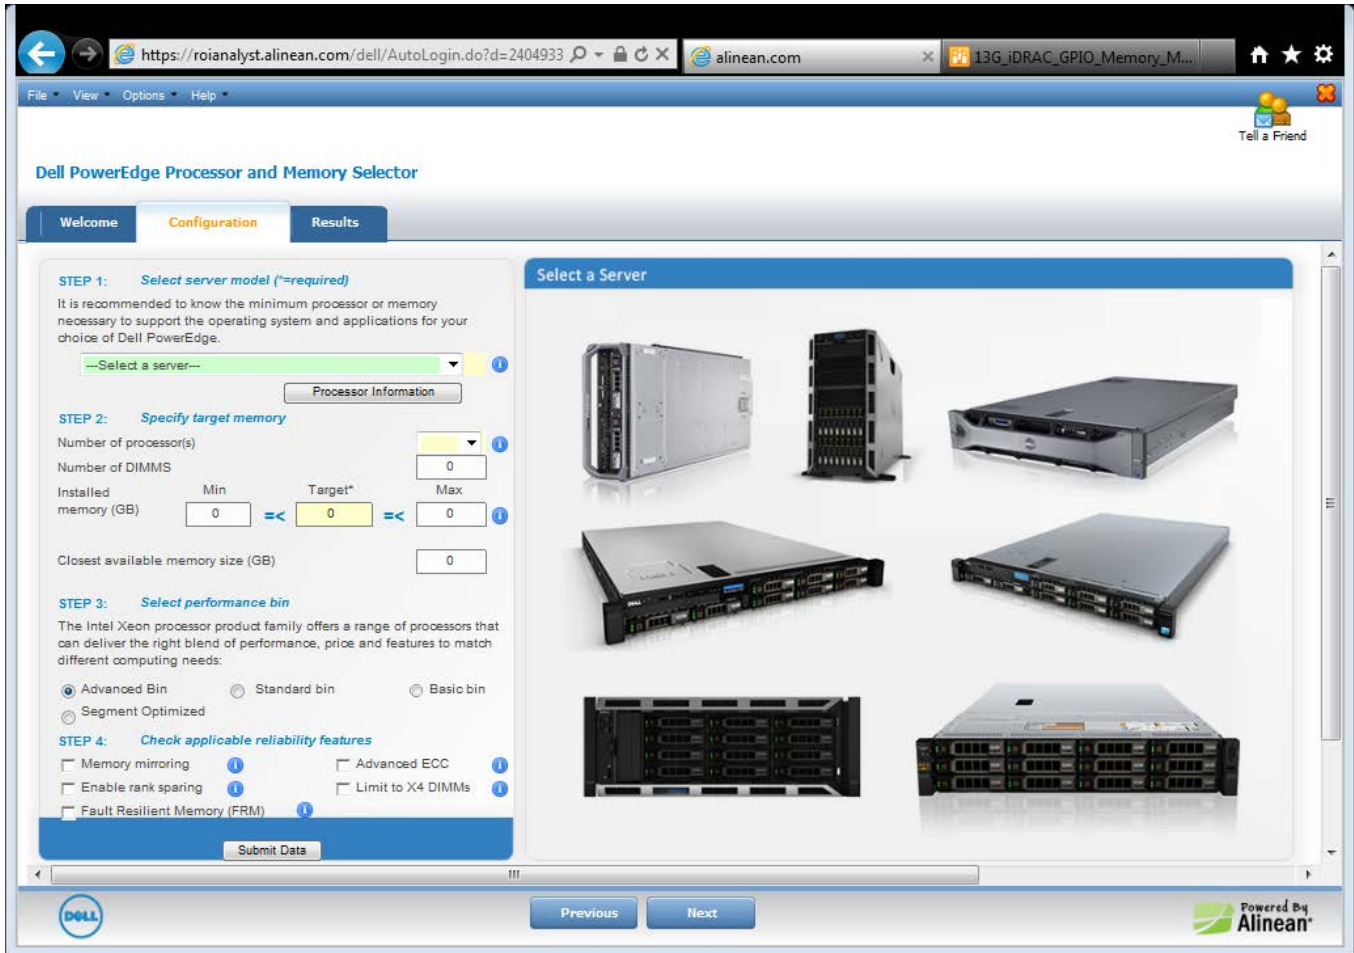


Figure 3 Dell PowerEdge Processor and Memory Selector - Configuration page





Once you've selected a server, the tool displays basic information about the server to help you decide which server most effectively meets your needs, as shown on the right panel of the Configuration page in Figure 4.

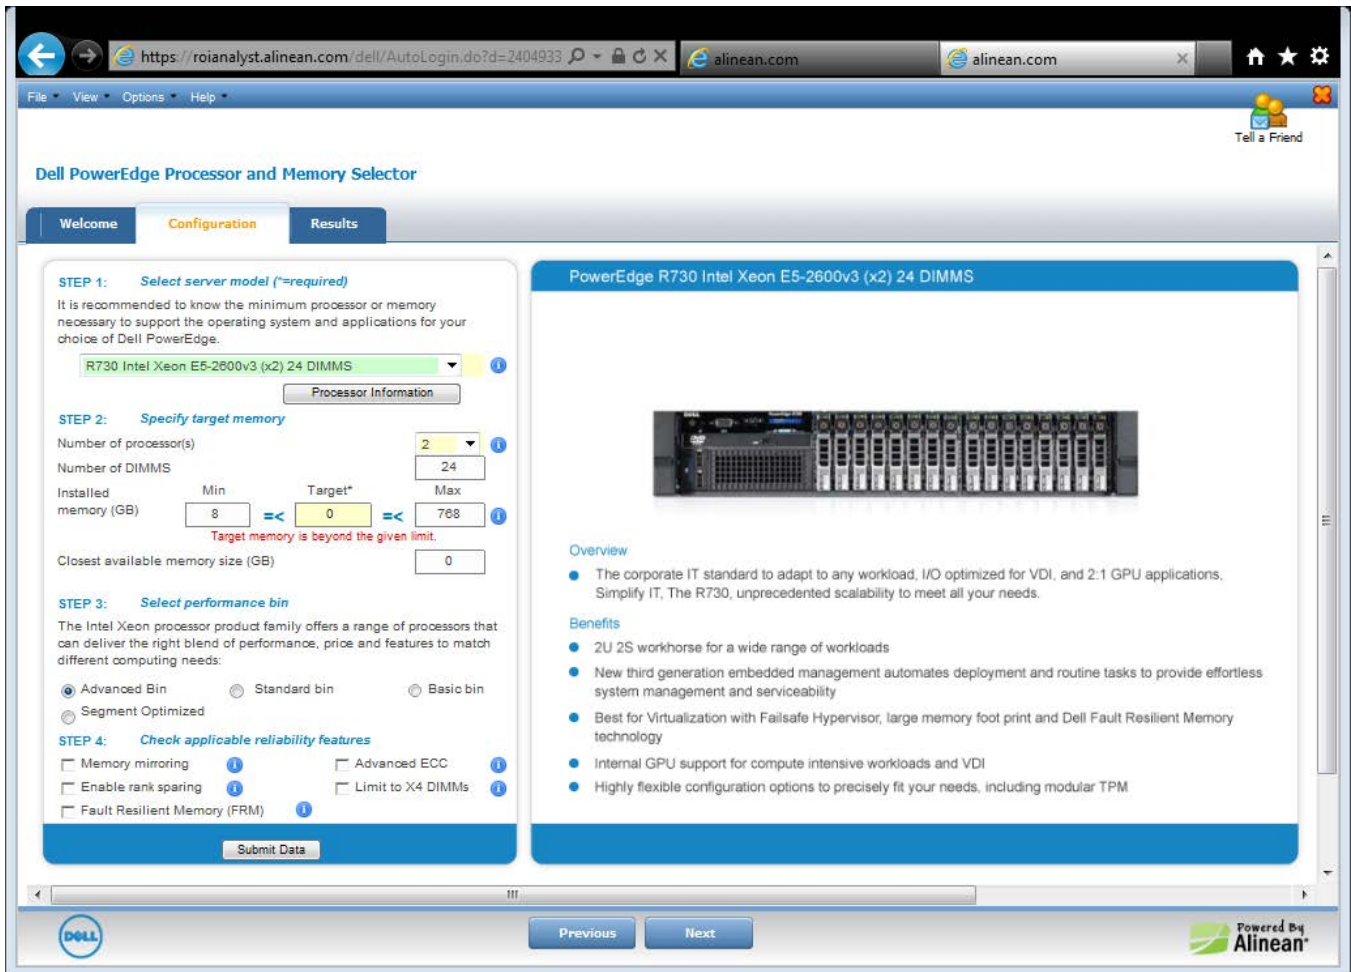


Figure 4 Dell PowerEdge Processor and Memory Selector - Server selection



Click **Processor Information** for detailed information on the processor family of the chosen server (Figure 5). On the Processor Information page, click the link (Figure 5, circled in red) for details about the processor product family.



Figure 5 Dell PowerEdge Processor and Memory Selector - Processor information

The processor page (Figure 6) provides detailed information about the platform-supported processor product family, including part numbers, core count, cache size, and clock frequency. The processors are divided into descriptive categories with application guidance for each category to help you match your requirements with the appropriate processor model.

**Intel® Xeon® processor E3-2600v3 product family**

The Intel® Xeon® processor E5-2600v3 product family has the performance, the efficiency and is focused on bring the Service Quality to enable IT managers to succeed.

**Selecting the Processor That Is Right for You**

The Intel® Xeon® processor E5-2600v3 product family offers a range of processors that can deliver the right blend of performance, price and features to match different computing needs:

- Advanced: Offers highest functionality and optimal server return on investment (ROI)
- Standard: Best mix of performance, value and advanced features
- Basic: Cost-sensitive customers seeking basic features
- Low power: Focused on energy efficiency and designed to deliver lowest power and best performance per watt

Frequency optimized: Ideal for high-performance computing (HPC) and technical computing with the highest frequency processor offering

	Processor Number	Cores/Cache	Frequency
Advanced	E5-2650 v3	10C/ 25M	2.3GHz
	E5-2660 v3	10C/ 25M	2.8GHz
	E5-2670 v3	12C/ 30M	2.3GHz
	E5-2680 v3	12C/ 30M	2.5GHz
	E5-2690 v3	12C/ 30M	2.6GHz
Basic	E5-2603 v3	6C/ 15M	1.8GHz
	E5-2609 v3	6C/ 15M	1.9GHz
Low Power	E5-2630L v3	8C/ 20M	1.8GHz
	E5-2650L v3	12C/ 30M	1.8GHz
Segment Optimized	E5-2687 v3	8C/ 20M	3.2GHz
	E5-2699 v3	18C/ 45M	2.3GHz
	E5-2623 v3	4C/ 10M	3.0GHz
	E5-2637 v3	4C/ 15M	3.5GHz
	E5-2643 v3	6C/ 20M	3.4GHz
	E5-2683 v3	14C/ 35M	2.0GHz
	E5-2685 v3 *	12C/ 30M	2.8GHz
	E5-2695 v3	14C/ 35M	2.3GHz
	E5-2697 v3	14C/ 35M	2.8GHz
	E5-2698 v3	16C/ 40M	2.3GHz
Standard	E5-2620 v3	6C/ 15M	2.4GHz
	E5-2630 v3	8C/ 20M	2.4GHz
	E5-2640 v3	8C/ 20M	2.6GHz
Workstation	E5-1620 v3	4C/ 10M	3.5GHz
	E5-1630 v3	4C/ 10M	3.7GHz
	E5-1650 v3	6C/ 15M	3.5GHz
	E5-1660 v3	8C/ 20M	3.0GHz
	E5-1680 v3	8C/ 20M	3.2GHz
	E5-2687W v3	10C/ 25M	3.1GHz

\* The E5-2685 v3 is an off-roadmap part & does not support VT or HT

Close

Figure 6 Dell PowerEdge Processor and Memory Selector - Processor details

After selecting a processor category, click **Close** to return to the Dell PowerEdge Processor and Memory Selector Configuration page (Figure 4).



Step 2 provides guidance for memory configuration, again based on user input. First, enter the number of processors for the server. Based on the number of processors, the **Installed memory (GB)** fields, comprising the maximum number of DIMMs and the minimum and maximum memory sizes, update automatically (Figure 7).

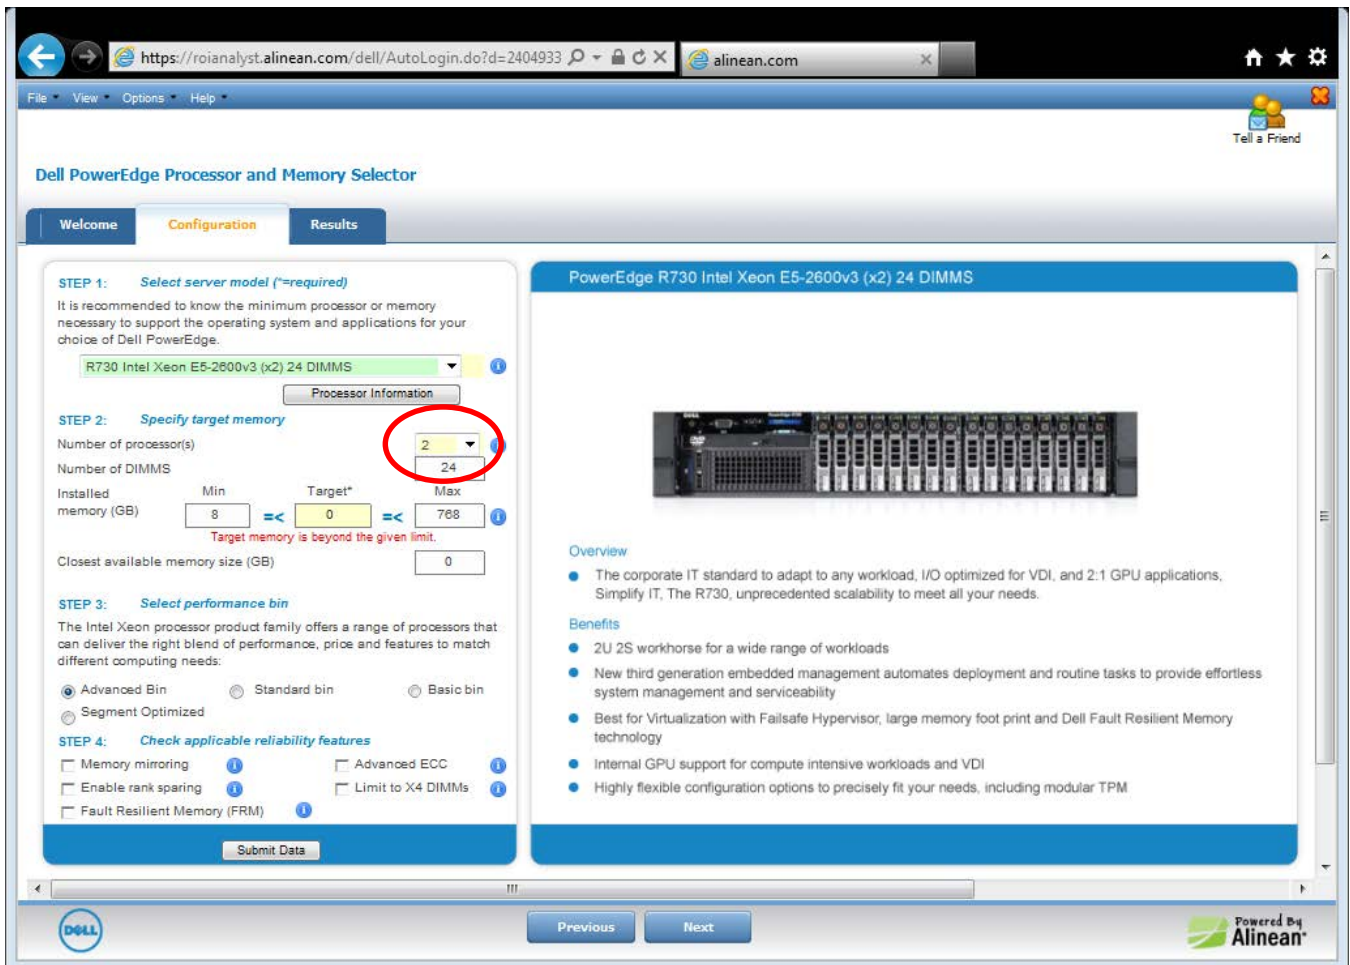


Figure 7 Dell PowerEdge Processor and Memory Selector – Setting processor count



Next, set the desired target memory size based on the server operating system, application, and workload (Figure 8). The target memory size must be within the range specified by the **Min** and **Max** fields (inclusive). The application automatically finds the closest available memory size matching the target memory value supplied.

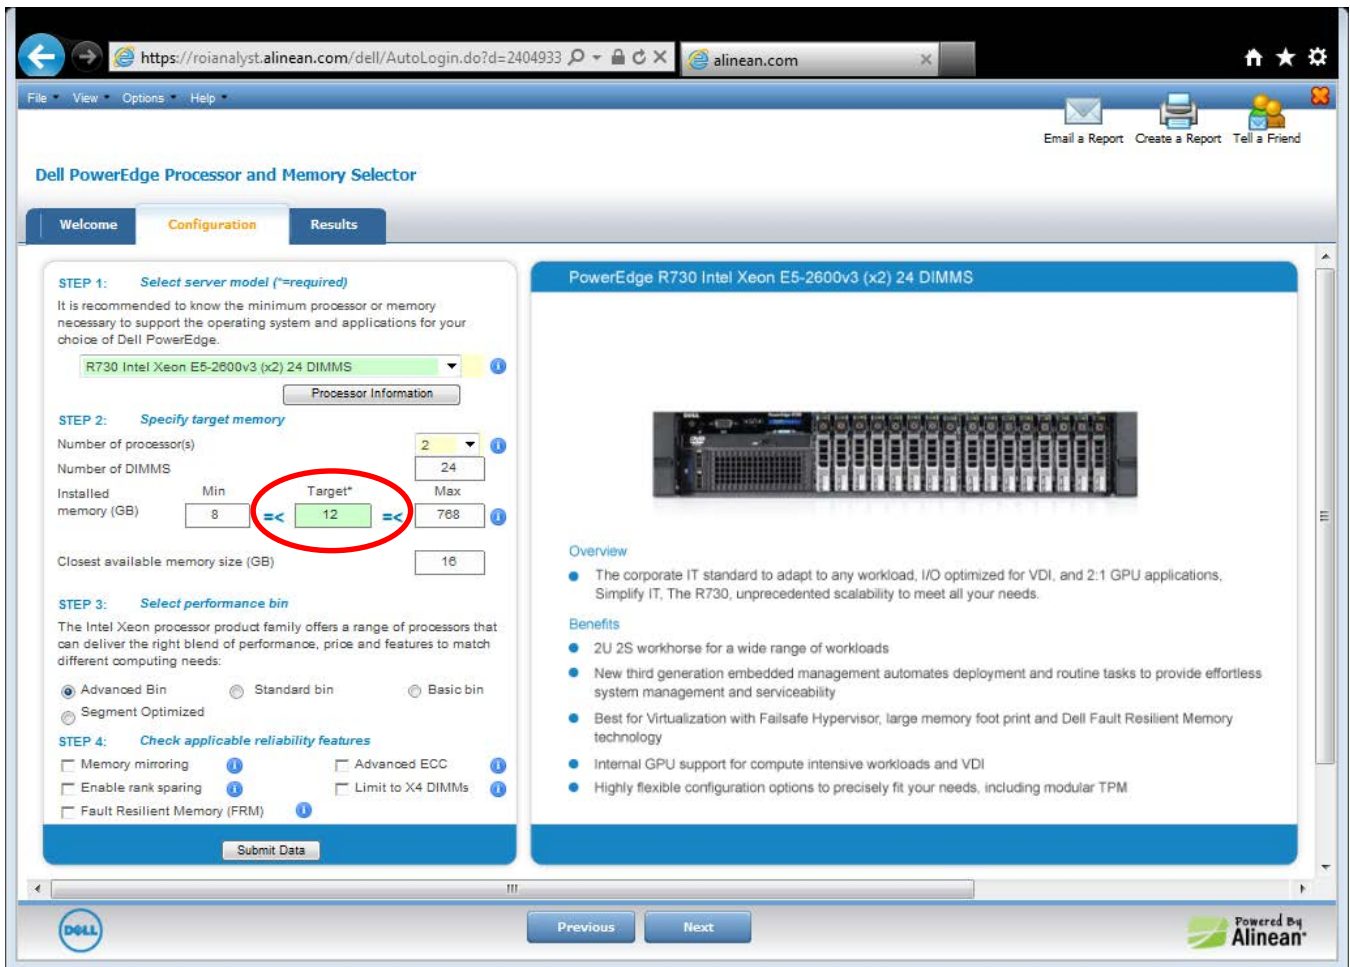


Figure 8 Dell PowerEdge Processor and Memory Selector – Setting target memory size





In Step 3, select the processor performance bin (Figure 9) based upon the information garnered from the Processor Information page (Figure 6) in Step 1.

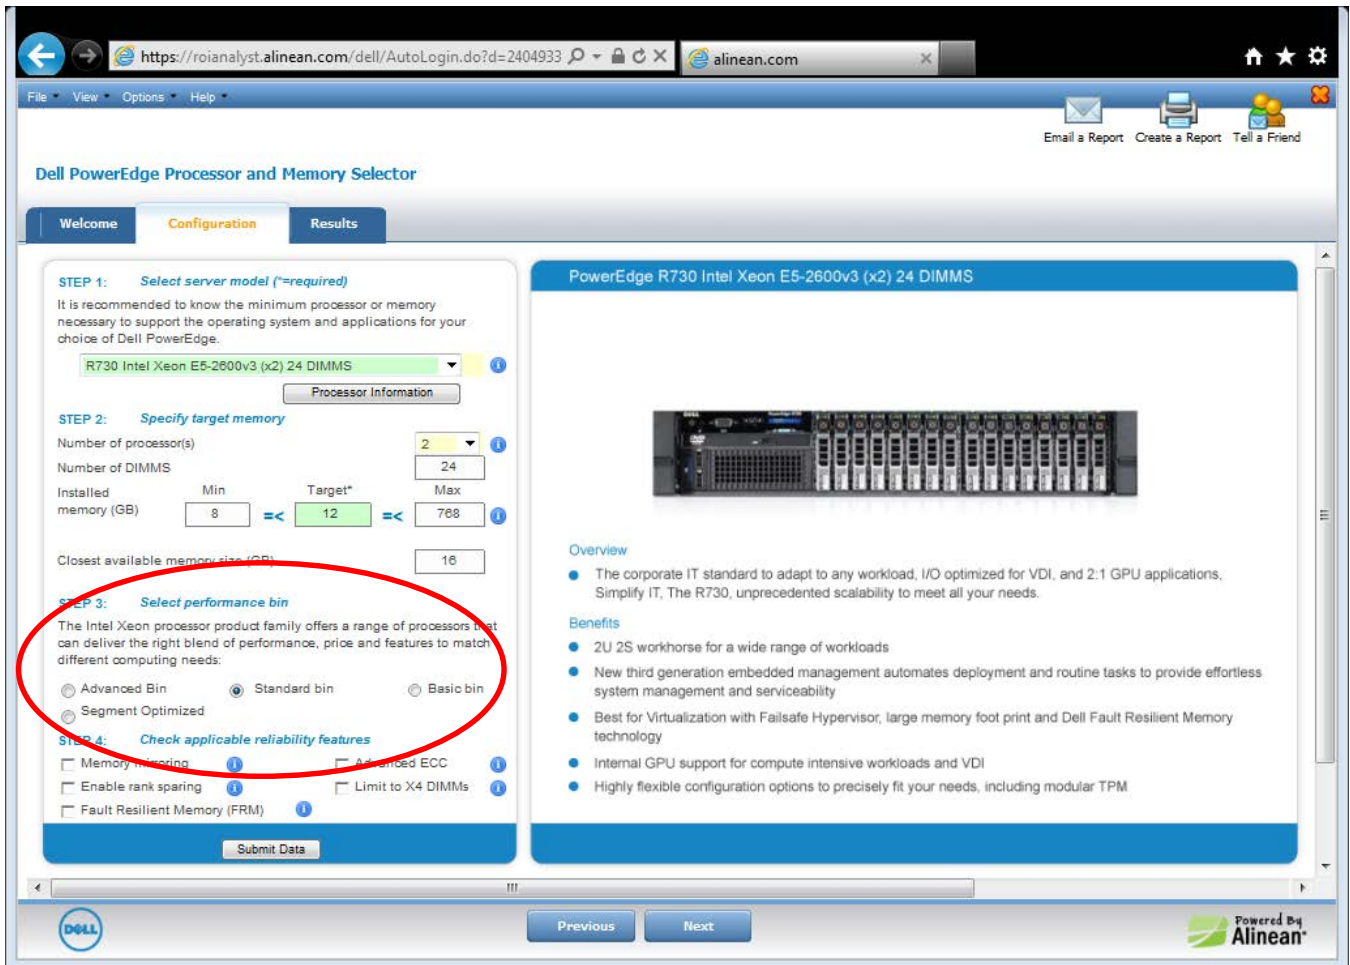


Figure 9 Dell PowerEdge Processor and Memory Selector – Setting processor performance bin



In Step 4, select the memory reliability features (Figure 10). The various choices here affect the final number of required DIMMs, and you may need to adjust the target memory size in Step 2.

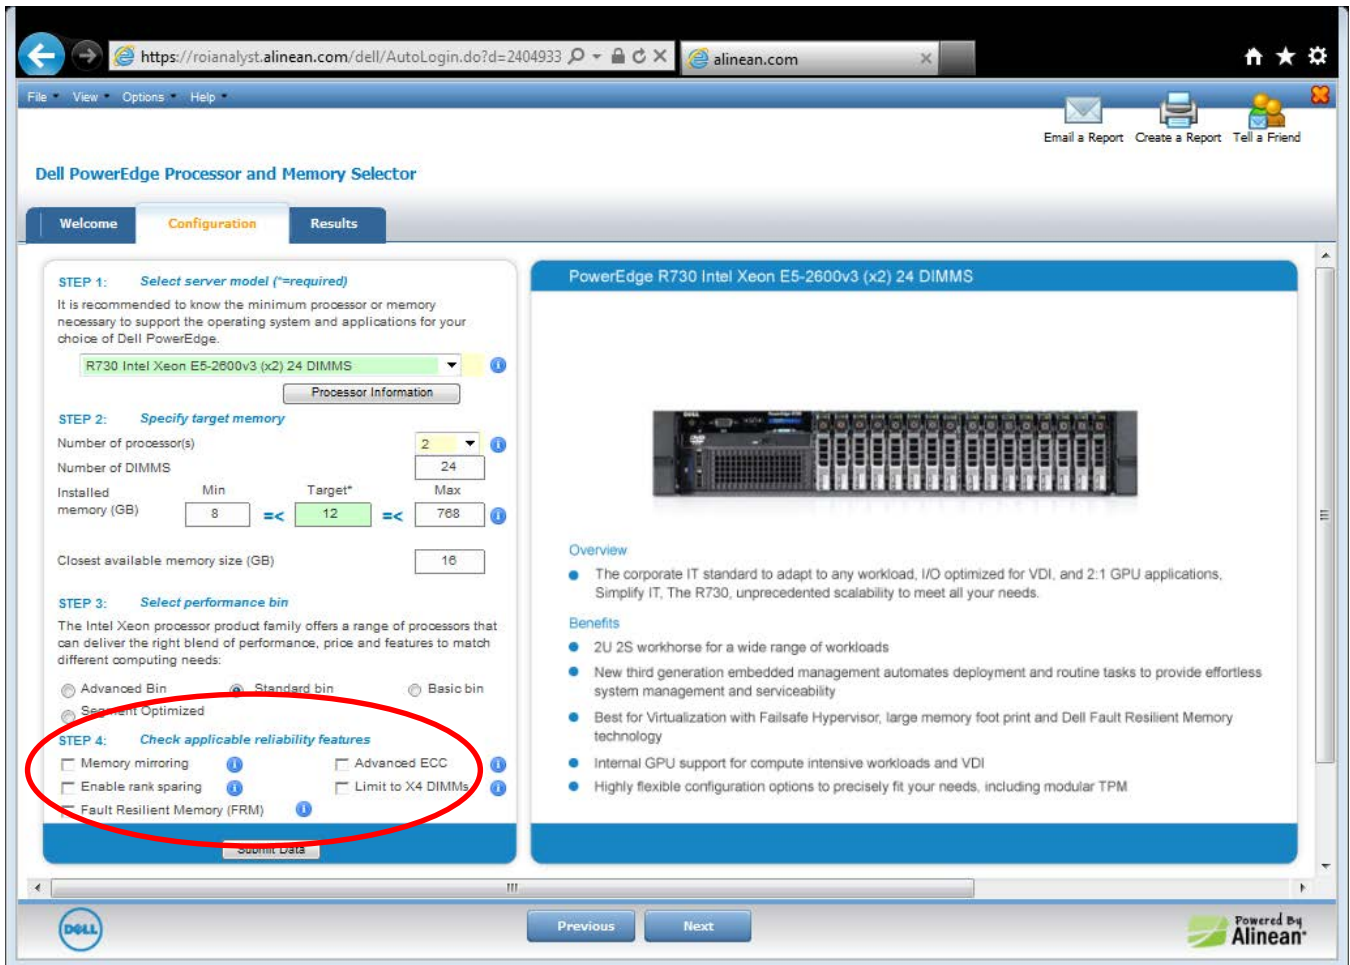


Figure 10 Dell PowerEdge Processor and Memory Selector – Memory reliability features



In the example shown in Figure 11, selecting the memory reliability option **Advanced ECC** causes the minimum limit of the installed memory to increase from 8GB to 16GB, which results in the requested value "Target" to fall outside of the allowed range of values.

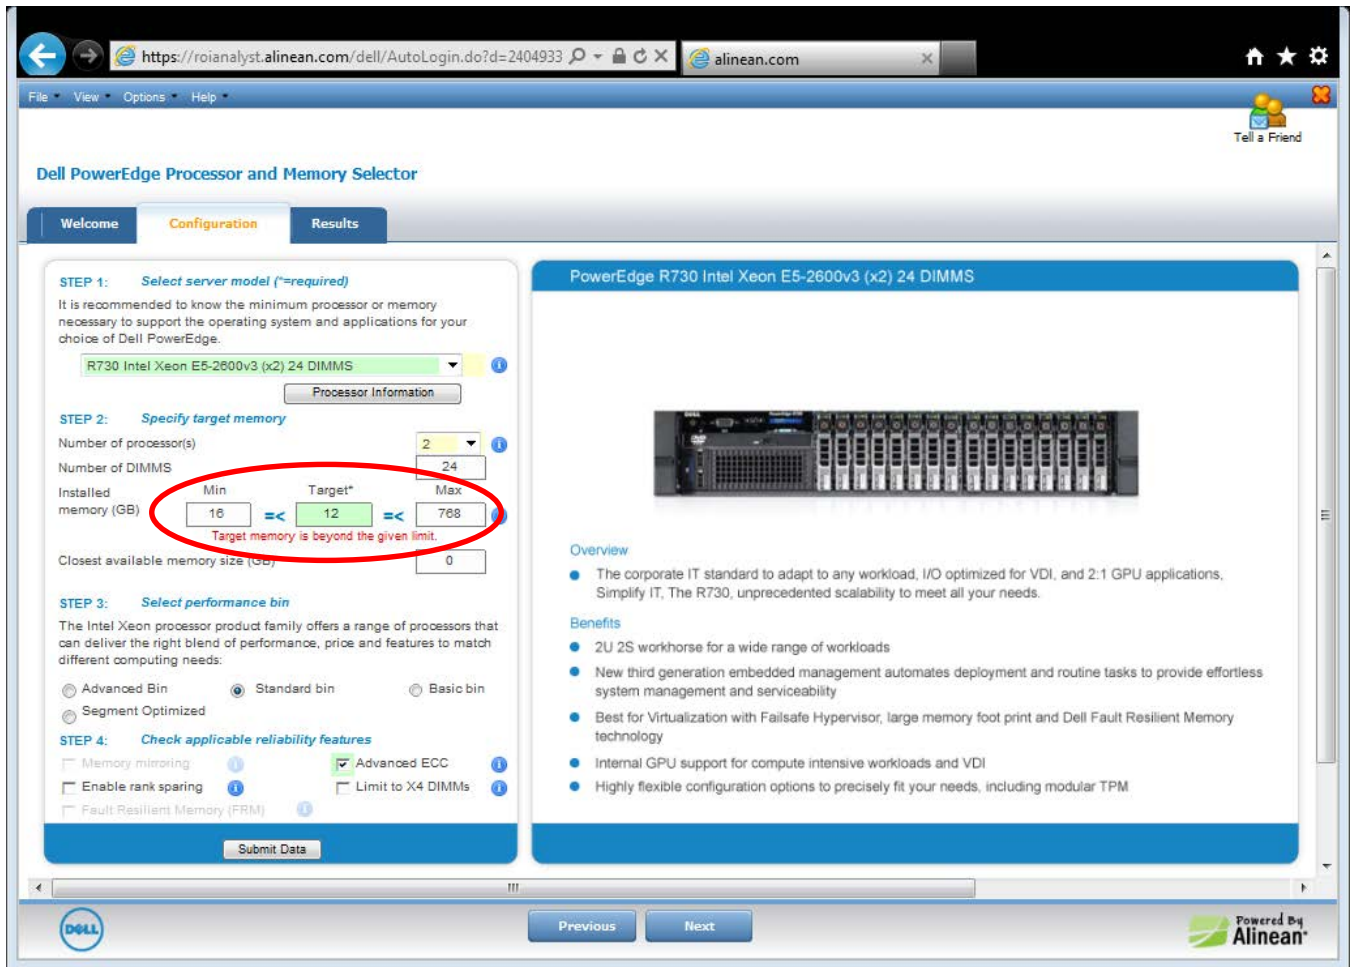


Figure 11 Dell PowerEdge Processor and Memory Selector – Target memory size invalid





Adjusting the target memory size to 16GB in the example corrects the error (Figure 12).

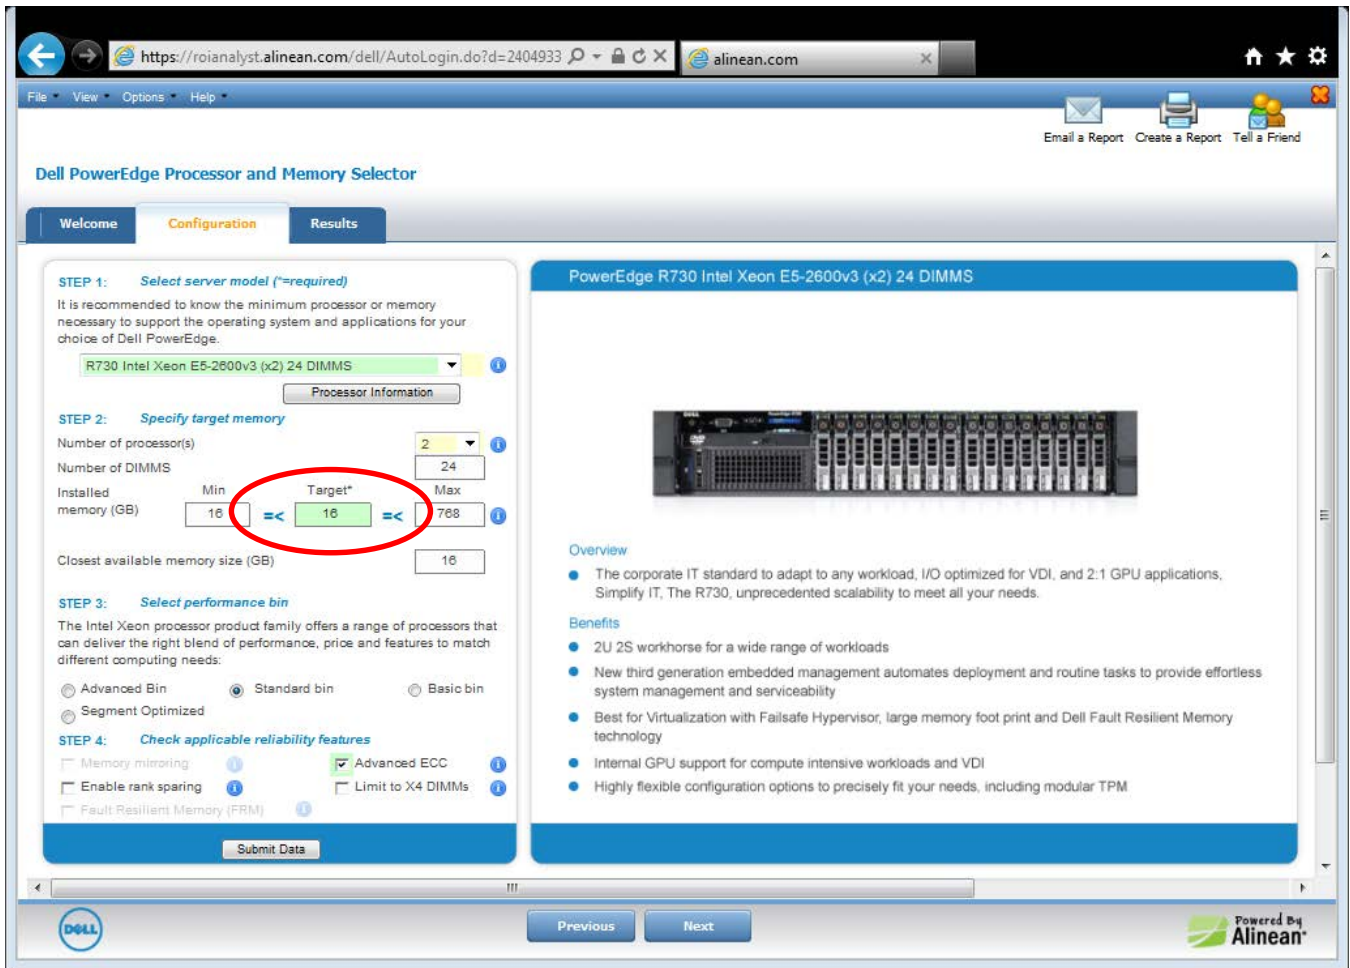


Figure 12 Dell PowerEdge Processor and Memory Selector – Target memory size correction



Once you are satisfied with the server configuration, click **Submit Data** or click **Next** to see the recommended results (Figure 13).

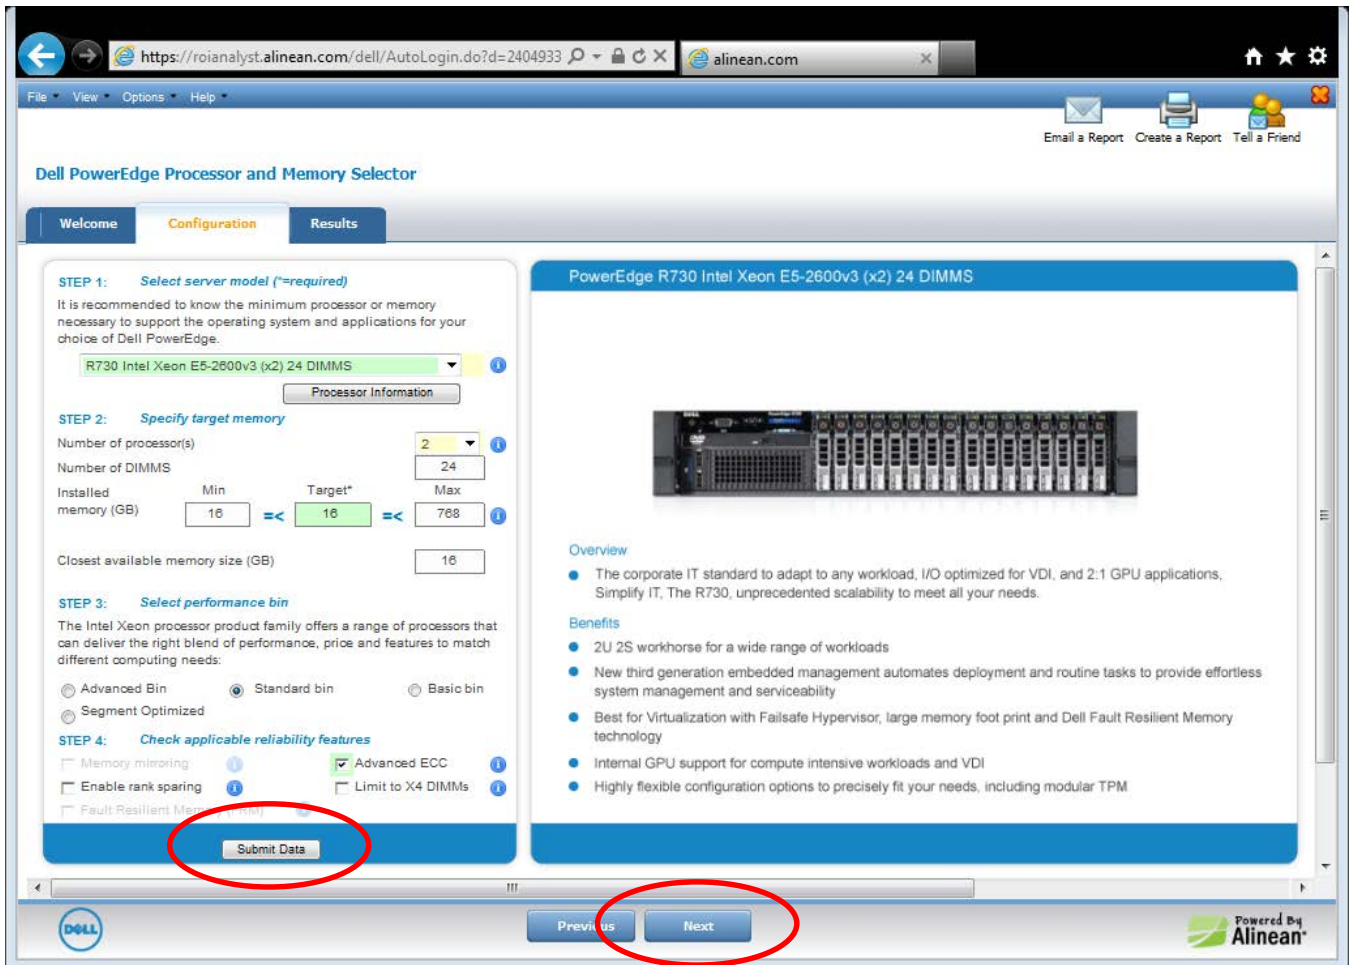


Figure 13 Dell PowerEdge Processor and Memory Selector – Submit data



In the example in Figure 14, the tool notifies the user that the chosen memory configuration does not use all of the available memory channels, creating an unbalanced configuration. To change the memory configuration, return to the Configuration page by clicking the Configuration tab or clicking Previous.

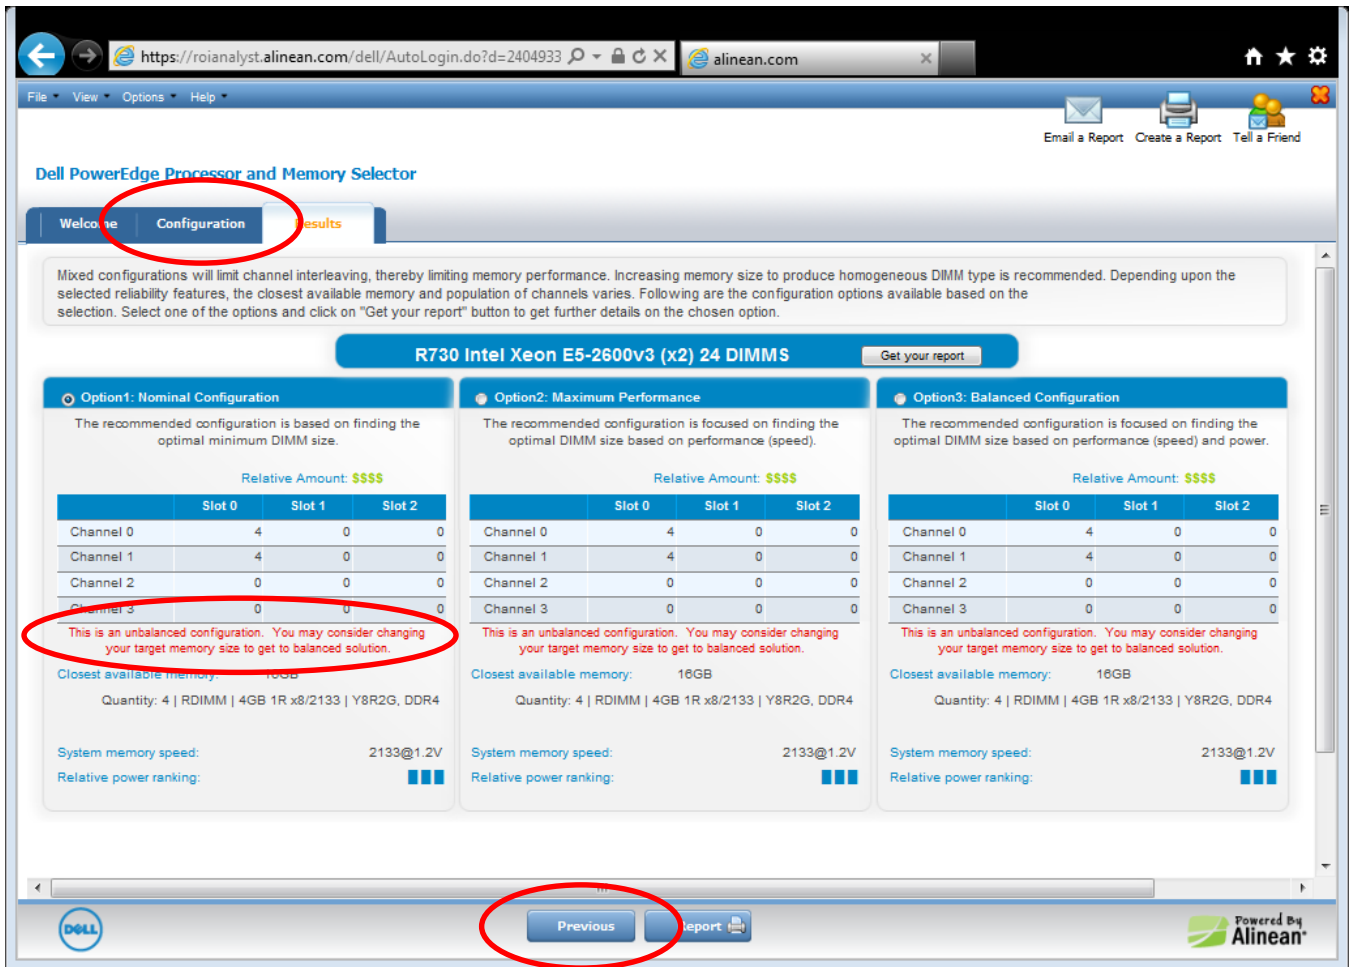


Figure 14 Dell PowerEdge Processor and Memory Selector – Results page for unbalanced configuration



The Dell PowerEdge Processor and Memory Selector Results page (Figure 15) displays multiple configurations that balance power consumption with system performance. In this example:

- **Option 1: Nominal Configuration** stresses finding the minimum DIMM size that meets the user requirements.
- **Option 2: Maximum Performance Configuration** finds the optimal DIMM size to maximize performance.
- **Option 3: Balanced Configuration** provides a memory configuration that provides a balance between maximizing performance and minimizing power consumption.

Mixed configurations will limit channel interleaving, thereby limiting memory performance. Increasing memory size to produce homogeneous DIMM type is recommended. Depending upon the selected reliability features, the closest available memory and population of channels varies. Following are the configuration options available based on the selection. Select one of the options and click on "Get your report" button to get further details on the chosen option.

**R730 Intel Xeon E5-2600v3 (x2) 24 DIMMS** [Get your report](#)

**Option1: Nominal Configuration**

The recommended configuration is based on finding the optimal minimum DIMM size.

Relative Amount: \$\$\$\$

	Slot 0	Slot 1	Slot 2
Channel 0	4	4	0
Channel 1	4	4	0
Channel 2	4	4	0
Channel 3	4	4	0

Closest available memory: 64GB  
Quantity: 16 | RDIMM | 4GB 1R x8/2133 | Y8R2G, DDR4

System memory speed: 1800@1.2V  
Relative power ranking: █

**Option2: Maximum Performance**

The recommended configuration is focused on finding the optimal DIMM size based on performance (speed).

Relative Amount: \$\$\$

	Slot 0	Slot 1	Slot 2
Channel 0	8	0	0
Channel 1	8	0	0
Channel 2	8	0	0
Channel 3	8	0	0

Closest available memory: 64GB  
Quantity: 8 | RDIMM | 8GB 2R x8/2133 | H8PGN, DDR4

System memory speed: 2133@1.2V  
Relative power ranking: █ █ █ █

**Option3: Balanced Configuration**

The recommended configuration is focused on finding the optimal DIMM size based on performance (speed) and power.

Relative Amount: \$\$\$

	Slot 0	Slot 1	Slot 2
Channel 0	8	0	0
Channel 1	8	0	0
Channel 2	8	0	0
Channel 3	8	0	0

Closest available memory: 64GB  
Quantity: 8 | RDIMM | 8GB 2R x8/2133 | H8PGN, DDR4

System memory speed: 2133@1.2V  
Relative power ranking: █ █ █ █

[Previous](#) [Report](#)

Powered By **Alinean**

Figure 15 Dell PowerEdge Processor and Memory Selector – Results page for final configuration



For a PDF-formatted report for each recommended configuration, click the radio button by the configuration option(s), and then click **Get your report** or click **Report** (Figure 16).

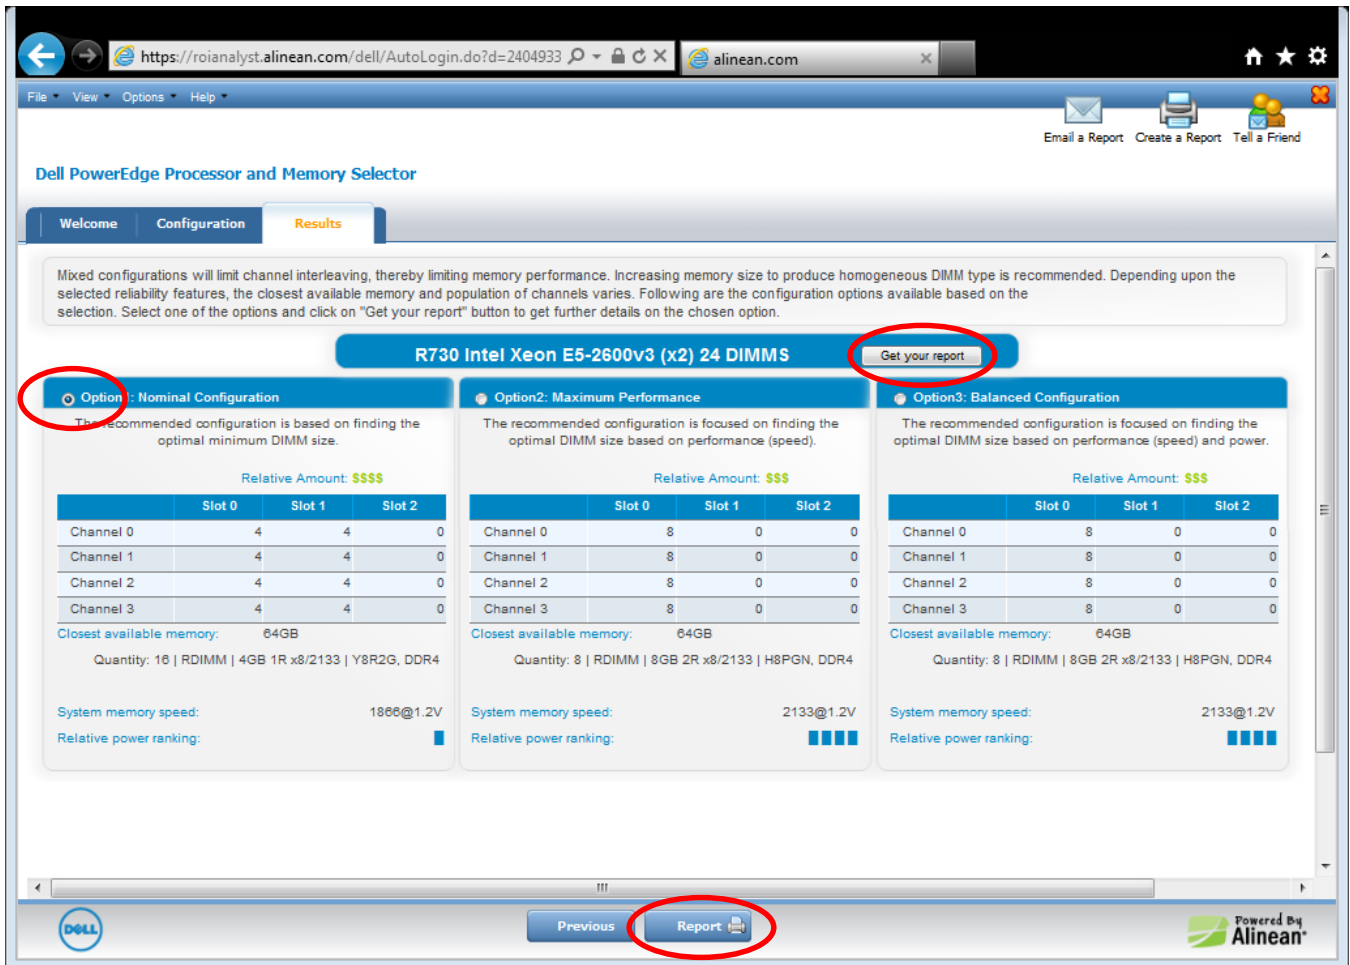


Figure 16 Dell PowerEdge Processor and Memory Selector – Get report



Select the “Dell PowerEdge Processor and Memory Selector Analysis” template and click **Create Report** (Figure 17).

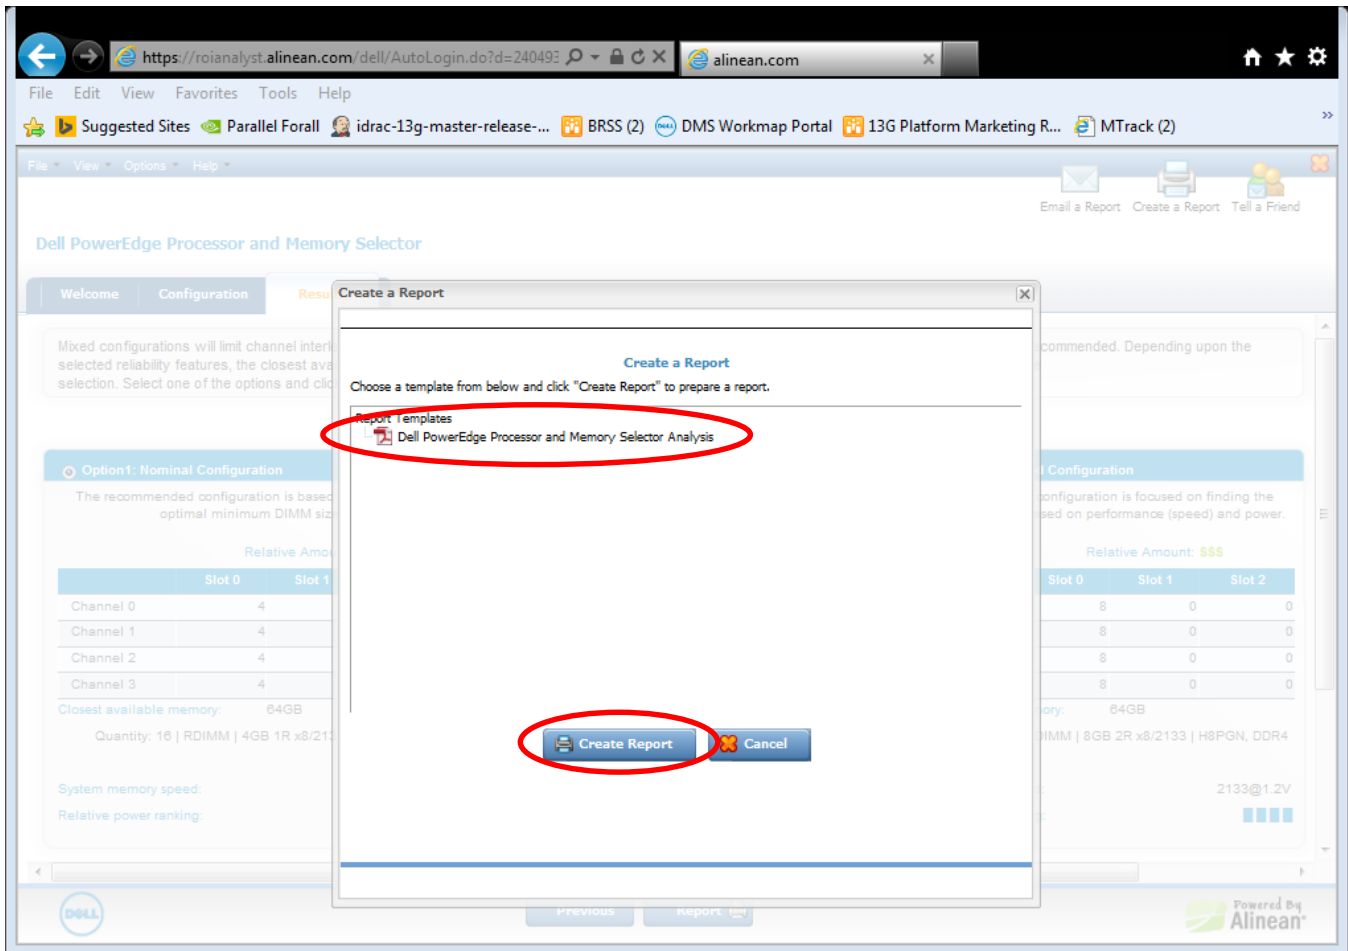


Figure 17 Dell PowerEdge Processor and Memory Selector – Report template

## Order validator

Dell ordering tools protect your Dell PowerEdge server orders by validating that the selected server configuration can be supported by the selected power supply(s). If the selected power supply cannot support the hardware configuration, the order validator directs you to upgrade the power supply or downgrade the server configuration, and blocks the order from being submitted. If for some reason you are unsuccessful in ordering a power supply configuration that the **Help Me Choose for PSUs** feature shows as supported, contact Dell Sales for assistance.




Figure 18 Order validator – Power supply error message



## Components

Various component options are available on Dell PowerEdge servers that can reduce power consumption. The overall power savings are compounded when power distribution (voltage regulator and power supply) efficiencies are also considered.

### Processor

For multiple generations, Dell PowerEdge servers have supported low-voltage versions of Intel Xeon processors that provide equivalent performance with lower power consumption. If one of the low-voltage processor options meets your performance and feature requirements, Dell recommends selecting the low-voltage processor version for additional power reduction.

### Memory

Many factors impact memory power consumption, including the DIMM type, organization, DRAM technology, DIMM capacity, operating frequency, and operating voltage. In addition to using the Memory Advisor tool, here are some general guidelines for selecting a memory configuration:

- **Consider higher capacity DIMMs to reduce the number of total DIMMs and DIMMs per channel (DPC).** As a general rule, two DIMMs of the same frequency and technology type will use more power than a single DIMM of twice the capacity. For instance, two 4GB DIMMs will use more power than a single 8GB DIMM. An additional benefit may be seen as reducing the DPC may allow a memory configuration to operate at a higher frequency and improve overall performance.
- **Select DIMMs with x8 DRAM.** DIMMs with x8 DRAM have half as many DRAM devices accessed at a time compared to DIMMs with x4 DRAM, reducing average power consumption.

### Storage

When selecting storage, there are many options that impact storage power consumption. Here are some general guidelines for selecting a storage configuration:

- **Consider higher capacity hard drives to reduce the number of total hard drives.** Workload should be carefully considered, as reducing the number of hard drives may impact the performance workloads sensitive to Input/Output Per Second (IOPS).
- **Consider solid-state drives (SSD).** While SSDs have limited capacity, they typically consume less power than hard disk drives that require a motor to spin the disk. They also have a much lower latency and larger bandwidth.
- **Consider lower RPM hard drives.** Typically, the lower the hard drive RPM, the lower the power consumption. Customer usage models that do not require top-end storage performance can benefit from lower power consumption.
- **Consider 2.5-inch hard drives.** Typically, hard drive power consumption is lower with smaller form factors.
- **Consider software RAID for configurations with four drives or less.** Software RAID is supported across the new PowerEdge server portfolio for configurations with four drives or less, which significantly reduces power consumption by removing a hardware RAID adapter from the configuration. While hardware RAID provides higher performance, configurations with four drives or less typically do not require top-end storage performance.





## Power supply

Power supply selection can have a significant impact on the power consumption of the system. The new PowerEdge server portfolio supports a broad power supply portfolio to help customers select a power supply appropriate for their hardware configuration. Power supplies are typically more efficient when operating at a larger percentage of the rated output. Selecting a lower rated output power supply (right-sizing) increases the percentage of the load to rated output. This is especially true for redundant PSU configurations that load share. When right-sizing, you should consider future server hardware upgrades to avoid the need for upgrading the power supply.

In addition to right-sizing, the power supply portfolio offers upgrade options for higher efficiency levels. When combined with right-sizing, the highest-efficiency power supply configuration can be selected for your configuration.

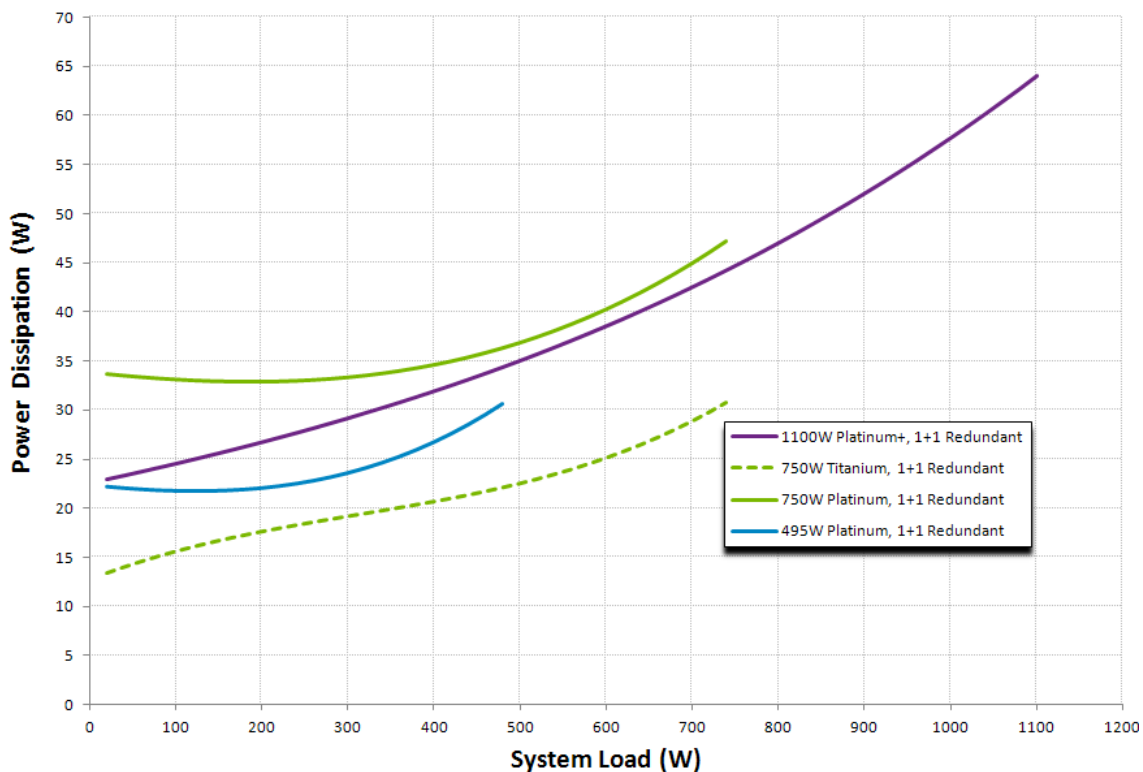


Figure 19 PSU power dissipation vs. system load





## BIOS Settings

The Dell PowerEdge 13<sup>th</sup> generation servers support system profiles that serve as “easy buttons” for applying default BIOS and firmware values targeted for performance, performance-per-watt, or RAS for dense configurations. You can select System Profiles as part of your order to be programmed in the Dell factory. If you are interested in power efficiency, a performance-per-watt profile is recommended. Dell specifically recommends the Dell Active Power Controller (DAPC) BIOS Setting in the ordering tool.

## How to configure a server for power efficiency

When you purchase a new PowerEdge server, you will receive a server that has been designed with a heavy design focus on power efficiency. However, there are many tweaks you can make to your server to save even more power. These additional adjustments take advantage of your particular usage model.

### No performance impact

Some tweaks may impact performance while others may not. The tweaks in this section do not impact performance.

### Hardware setup

C-states are low-power processor states that help with IDLE power. USB devices like mice and keyboards can impact processor C-states by preventing the system from entering low-power states. USB devices also require power from the server, increasing the server’s total power consumption. Using remote access and removing USB devices that are not required can often reduce the server’s power consumption.

### Unused services and ports

System BIOS provides multiple setup options to disable unused devices and ports. Disabling these devices allows the BIOS to configure them to reduce consumed power. Table 1 lists the BIOS setup options.

Table 1 Setup options: Unused devices and ports

Location	Submenu	Setting	Recommendation
System BIOS Settings	SATA settings	Embedded SATA	Disable if not used to reduce power.
System BIOS Settings	Integrated devices	Integrated RAID controller	Disable if not used to reduce power.
System BIOS Settings	Integrated devices	User-accessible USB ports	Disable if not used to reduce power.
System BIOS Settings	Integrated devices	Internal USB port	Disable if not used to reduce power.
System BIOS Settings	Integrated devices	Integrated network card 1	Disable if not used to reduce power.
System BIOS Settings	Integrated devices	Slot disablement	Disable if not used to reduce power.
iDRAC Settings	Network	Enable NIC	Disable if not used to reduce power.

### Power supply

Power supplies are optimized to run with high-line (230V) input voltage. While PowerEdge power supplies support auto-ranging, low-line (115V) operation can reduce overall efficiency by about 2%.

When you order redundant power supply configurations, overall system power efficiency is reduced by power supplies sharing the load and operating at a lower point on their efficiency curves. Rack and tower PowerEdge



servers support Power Supply Hot Spare, which puts the redundant power supply in a “sleep state” to remove the efficiency impact of redundant power supply configurations. Hot Spare is enabled by default, and you can select which of the power supplies will be the hot spare primary. This should be the power supply on the most cost-effective and power-efficient grid, as the primary power supply may see the majority of the load. Rack and tower PowerEdge servers also support the ability to disable power factor correction (PFC) to reduce power consumption when the system is in standby (S5 state).

Table 2 lists the setup options for configuring the power supply.

Table 2 Setup options: Power supply

Location	Submenu	Recommendation
Power Configuration	Power supply redundancy policy	Set to redundant if PSU can support server hardware configuration.
Power Configuration	Enable hot spare	Redundant PSU configuration required.
Power Configuration	Primary power supply unit	Select PSU connected to most efficient or cost-effective power grid.
Power Configuration	Enable power factor correction	Disable to reduce S5 standby power.

## Component updates

Dell periodically releases updates for components such as BIOS, drives, and firmware. These updates may provide improvements to performance and power consumption. Install the latest updates to ensure your servers have access to the latest improvements.

## Potential for minor performance impact

New PowerEdge servers support multiple tweaks that may have a minor impact to performance. The performance impact is often workload-dependent and may not be measurable with typical workloads.

## System profiles

As first described in “How to order a server”, system profiles serve as a quick and easy way to apply default BIOS and firmware values targeted for performance, performance-per-watt, or RAS for dense configurations. If you are interested in power efficiency, a performance-per-watt profile is recommended. Dell specifically recommends the Performance per Watt (DAPC) profile. This is the system’s default setting.

## Performance per Watt (DAPC)

DAPC is a Dell proprietary implementation of dynamic processor power management that provides superior performance-per-watt compared to OS-based dynamic power management. Unlike unique OS implementations, DAPC provides consistent results regardless of the operating system and provides dynamic processor power management even if it is not supported by the OS or hypervisor.

DAPC System profile is configured in the factory when you select “Power Saving Dell Active Power Controller” in the ordering tool. This is the recommended profile for performance-per-watt.



## Performance per Watt (OS)

This profile uses OS-based dynamic processor power management. The effectiveness will vary based on the operating system implementation.

The Performance per Watt (OS) System profile is configured in the factory when you select the “Power Saving BIOS Setting” option in the ordering tool.

Table 3 provides the setup options for configuring the system profile.

Table 3 Setup options: System profile

Location	Submenu	Setting	Description
System BIOS Settings	System Profile settings	System Profile	Select Performance per Watt (DAPC) profile.

Table 4 provides the setup option values associated with each system profile.

Table 4 System profile definitions

Subsystem System	Setting	System Profile				
		Performance	Performance per Watt (OS)	Performance per Watt (DAPC) (default)	Dense	Custom
CPU	CPU Power Management	Maximum Performance	OS DBPM	System DBPM (DAPC)	System DBPM (DAPC)	System DBPM (DAPC)
	Turbo Boost	Enabled	Enabled	Enabled	Disabled	Enabled
	C1E	Disabled	Enabled	Enabled	Enabled	Enabled
	C-States	Disabled	Enabled	Enabled	Enabled	Enabled
	Monitor/MWait	Enabled	Enabled	Enabled	Enabled	Enabled
Memory	Frequency	Maximum Performance	Maximum Performance	Maximum Performance	Dense Configuration Optimized	Maximum Performance
	Patrol Scrub	Standard	Standard	Standard	Extended	Standard
	Refresh Rate	1x	1x	1x	2x	1x
Thermal	Thermal Algorithm	Maximum Performance	Minimum Power	Minimum Power	Dense Configuration Optimized	Maximum Performance

## Power management

This section refers to the dynamic management of a component’s performance for a given workload to reduce power consumption.

Processor power management controls the use of processor performance states (P-states), where the processor operating frequency and voltage are dynamically adjusted to match workload performance requirements. For typical workloads, demand-based power management (DBPM) is recommended, particularly System DBPM (DAPC). Latency sensitive and deterministic-performance workloads could be



impacted by the delays required to transition between P-states. To disable P-state transitions and force the processor to always run at the rated frequency, set to Maximum Performance.

Processor C-states are IDLE, non-operational, low-power processor states. While C-states reduce IDLE power consumption, latency sensitive and deterministic-performance workloads could be impacted by the delay to exit a C-state. Deeper C-states provide greater power savings but have longer exit delays.

Memory Clock Enable (CKE) and self-refresh power management features reduce memory power when IDLE. Like processor P- and C-states, latency sensitive and deterministic-performance workloads could be impacted. PowerEdge servers also support power management of the Quick Path Interconnect (QPI) interface on Intel Xeon processors. QPI power management also introduces latency penalties for returning to a performance state. Memory and QPI power management are tied to the C1E and C-state setup options and can be disabled.

Dell PowerEdge RAID controllers (PERC) support Physical Disk Power Management through the Dell OpenManage Storage Management application. Like power management features for processor and memory, Physical Disk Power Management spins down IDLE drives to reduce power consumption. Physical Disk Power Management is disabled by default.

Table 5 provides the setup options for configuring power management.

Table 5 Setup options: Power management

Location	Submenu	Setting	Description
System BIOS Settings	System Profile Settings	CPU Power Management	Enable System DBPM (DAPC).
System BIOS Settings	System Profile Settings	C1E	Enable C1E. Also enables DIMM CKE power down and QPI power management.
System BIOS Settings	System Profile Settings	C States	Enable C-states. Also enables DIMM self-refresh.
System BIOS Settings	System Profile Settings	Monitor / Mwait	Disable for rogue OS that does not comply with C-states disabled. Can be set to disabled only when the C-state option is set to disabled.
Open Manage Server Administrator	Storage	Power Saving Mode	Maximum Power Saving Mode to reduce IDLE power of all drives.
Open Manage Server Administrator	Storage	Time Interval for Spin Down (in Minutes)	Shorten time for more aggressive power savings.
Open Manage Server Administrator	Storage	Spin Down Unconfigured	Enable to spin down unconfigured drives to reduce power.
Open Manage Server Administrator	Storage	Spin Down Hot Spares	Enable to spin down Hot Spare drives to reduce power.
Open Manage Server Administrator	Storage	Spin Down Configured	Enable to spin down configured drives to reduce IDLE power.



Location	Submenu	Setting	Description
Open Manage Server Administrator	Storage Saving (Idle C)	Automatic Disk Power	Enable Idle C to reduce IDLE power.

## Reduce top-end power consumption

You can also reduce power by limiting top-end performance. The performance impact of reducing top-end performance is workload dependent, and may be acceptable for many workloads. Turbo Boost is a performance feature in Intel processors that provides temporary operation above the rated processor frequency. The benefit of Turbo Boost is based on power and thermal headroom where there can be performance variations; for example, between processor sockets and servers. Workloads requiring deterministic performance may need to disable Turbo Boost.

Turbo Boost does present some challenges to the system due to the additional power required to support the performance bursts. The additional power allocation that is required to run in Turbo Boost mode can impact power supply right-sizing. If Turbo Boost is disabled, its power allocation can be removed from the platform power budget to support a smaller power supply, or it can be re-allocated to support additional components, such as memory or storage.

Reducing memory frequency is another option to reduce power by reducing top-end performance. Reducing memory frequency will have a direct impact on DIMM power consumption.

Table 6 highlights setup options to reduce top-end power consumption.

Table 6 Reducing top-end power consumption

Location	Submenu	Setting	Description
System BIOS Settings	System Profile Settings	Turbo Boost	Disable to reduce peak power. Requires Custom System Power.
System BIOS Settings	System Profile Settings	Memory Frequency	Set to lower frequency to reduce power. Requires Custom System Power.
Processor Settings	QPI Speed	QPI Speed	Set to lower speed to reduce power.
System BIOS Settings	Processor Settings	Number of Cores per Processor	Minimize active physical cores to reduce power.
Devices	Device Specific	Link Speed	For each adapter, review settings for power options. For example, set to lower speeds to reduce power.

## Thermal management

Thermal management options can also impact power consumption. By default, the Thermal Base Algorithm is set to Auto, which maps to the selected system profile. If a Performance per Watt system profile was not selected, the Thermal Base Algorithm can be set to Minimum Power, which will reduce the fan speed response in high ambient environments, which will reduce total system power that may have slight performance impacts. Minimum Power setting provides a balance of performance and power and is the Thermal Base Algorithm setting mapped to the Performance per Watt system profiles.



Table 7 provides the setup options for thermal management.

Table 7 Setup options: Thermal management

Location	Submenu	Setting	Description
iDRAC Settings	Thermal	Thermal Base Algorithm	Set to Auto if System Profile is Performance per Watt (DAPC). Otherwise, set to Minimum Power.

## Miscellaneous

Table 8 describes miscellaneous power setup options.

Table 8 Miscellaneous power setup options

Location	Submenu	Setting	Description
System BIOS Settings	Processor Settings	Adjacent Cached Line Prefetcher	Disable for performance per watt.
System BIOS Settings	Processor Settings	Hardware Prefetcher	Disable for performance per watt.
System BIOS Settings	Processor Settings	DCU Streamer Prefetcher	Disable for performance per watt.
System BIOS Settings	Processor Settings	DCU IP Prefetcher	Disable for performance per watt.
System BIOS Settings	System Profile Settings	Memory Refresh Rate	Set to 1x to reduce power.
System BIOS Settings	System Security	AC Power Recovery Delay	Select random or specify to minimize datacenter power spikes after AC recovery.
System BIOS Settings	Miscellaneous Settings	In-System Characterization	Enable for optimal tuning for power. May increase boot time after configuration change.

## Power capping

Dell PowerEdge 13<sup>th</sup> generation servers support a user-defined power cap through iDRAC8 and OpenManage Power Center. If data center-level efficiency curves are understood, power capping can be used to limit server power consumption to optimize overall data center power efficiency. Power capping can also be used to reduce peak power consumption. Depending on workload, performance impact may be negligible. See Figure 19 for an example of power capping.

Table 9 provides the setup options for configuring power capping.

Table 9 Setup options: Power capping

Location	Submenu	Setting	Description
Power Configuration	iDRAC Power Cap Policy	Enable if system power limit is desired.	Set to limit sustained power consumption.
iDRAC Settings	Power Configuration	Maximum Power Limit (Watts, BTU/hr)	Set to limit sustained power consumption.
iDRAC Settings	Power Configuration	Maximum % of Recommended System Limit	Set to limit sustained power consumption.



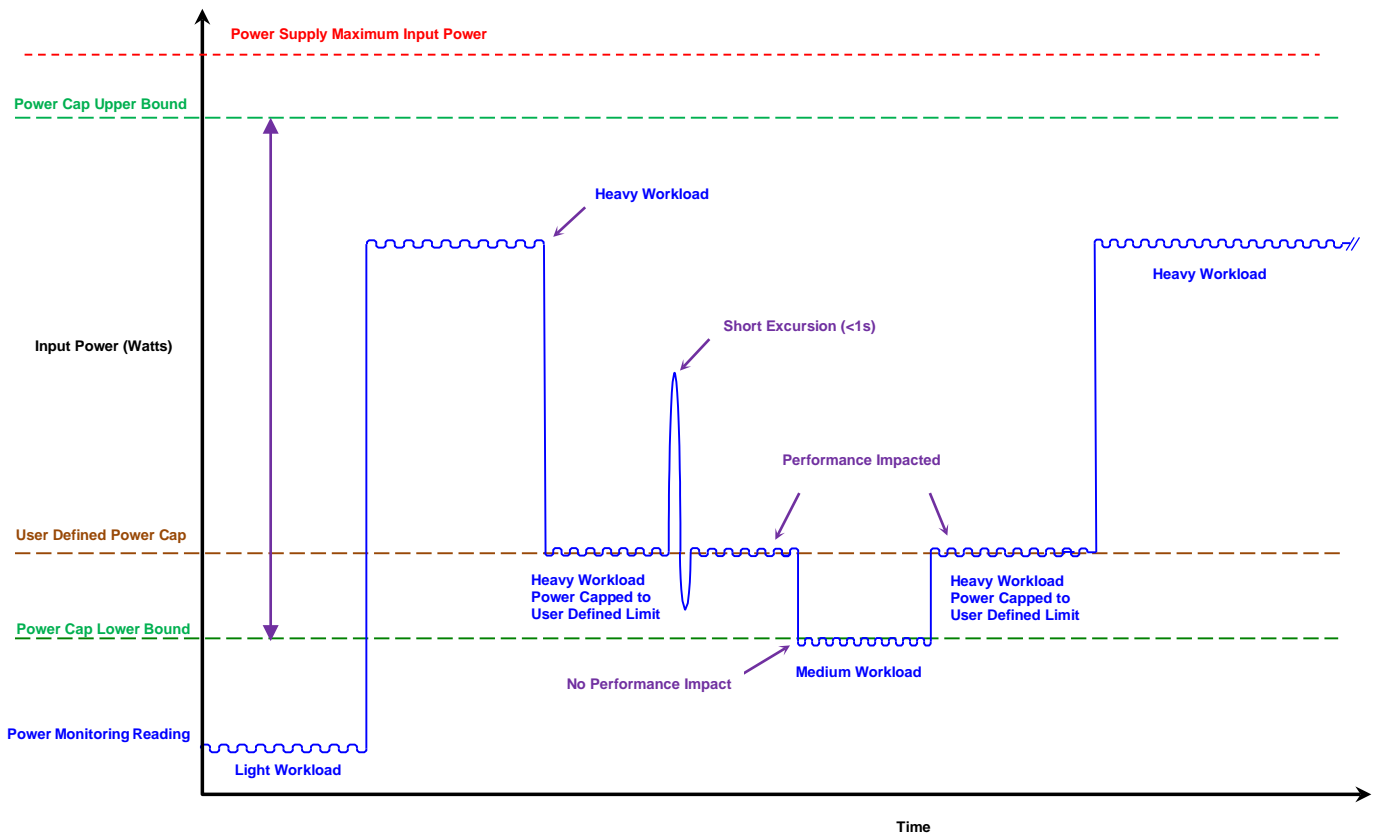


Figure 20 Power capping example

## Operating systems

Operating systems provide various power management configuration options such as USB selective suspend and video timeout. For details, review the operating system documentation on operating system power management features.

## About the authors

John Jenne is a Distinguished Engineer in the Dell Enterprise Product Group focused on Node Power across the Dell Server portfolio. John has previously held positions with Compaq, MaXXan Systems, and Newisys. He has a B.S. in Computer Engineering from Clemson University, an M.S.E.E. from the University of Houston, and currently holds 16 patents.

Craig Klein is a Principle Engineer in the Dell Enterprise Product Group working on Node Power across the Dell Server portfolio. Craig has previously held positions with IBM and Jabil Circuit. He has a B.S. in Electrical Engineering from Texas A&M University, and currently holds 148 patents.

