

Full Redundancy vs. Fault Tolerant Redundancy for PowerEdge Server PSUs

Tech Note by

Matt Ogle John Jenne David Hardy

Summary

Understanding the power supply redundancy options to facilitate your server is important for users seeking to prioritize certain use cases over others, such as full, consistent performance during fault conditions or higher performance and capabilities during normal operating conditions.

This DfD will discuss two PSU redundancy options; Full Redundancy (FR) and Fault Tolerant Redundancy (FTR), and explain when it may be advantageous for a user to adopt one of these solutions over the other.

Introduction

Customers need power redundancy to maintain application uptime. However, few know that there is more than one type of redundancy to consider, and the best option depends on several factors. This DfD will explain two power supply unit (PSU) redundancy options – **Full Redundancy (FR)** and **Fault Tolerance Redundancy (FTR)**. Dell Technologies now enables customers to select between these at point of sale for select platforms. Understanding these PSU redundancy options is critical as the selection will determine the minimum PSU capacity required to support the targeted PowerEdge server configuration.

FR configurations run at full performance during normal operating conditions and after PSU redundancy loss (if a PSU goes down due to input loss or fault). FR is optimized for consistent performance, thus the minimum PSU capacity allowed will ensure that the platform configurations full performance power requirements can be supported. **In summary** – PowerEdge users looking to adopt FR gain consistent PSU performance during normal and fault operating conditions, but will require a PSU capacity capable of supporting full performance power requirements.

FTR configurations run at full performance during normal operating conditions, but after PSU redundancy loss, intelligent platform power control loops may dynamically reduce system performance to limit the platform's power consumption within the capacity of the healthy PSU. FTR is optimized to enable support for richer platform configurations within a target PSU capacity that provides additional performance and capabilities during normal operations. The target PSU capacity is driven by multiple potential factors, such as:

- A larger PSU capacity is not available
- PSU capacity is right sized for a typical workload for CapEx and/or OpEx savings
- Require configuration support within the capacity of PSUs with C14 inlet connector
- Require configuration support within the low-line AC (110V) power limits of C14 and C20 inlet connectors
- Require PSU efficiency level and/or input type that is only available in limited PSU capacities

D&LLTechnologies

To support richer configurations with more perfomance and capability during normal operation, FTR takes advantage of the additional PSU capacity from the redundant PSU during normal operation. However, when the redundant PSU fails, FTR must take away performance to compensate for loss of additional power capacity that enabled the additional perfomance and capability. **In summary** – PowerEdge users looking to adopt FTR will have richer platform configuration options within a PSU capacity limit, but must assess the potential impact of performance degradation to their workload.

Addressing the Negative Stereotype

Historically, FR has been deemed as the superior PSU redundancy option. Customers viewed FTR concepts as a "trick" to compensate for a design limitation. Dell Technologies was originally opposed to supporting FTR due to the negative stigma associated with it.

Eventually, Dell Technologies added support for FTR to PowerEdge platforms because platform power requirements were increasing faster than PSU technology advancements. FTR was not advertised or marketed despite being an essential technology to support platform configurations that customers wanted. Only limited references were made in technical white papers.

As FTR concepts have become standard within the industry, it is now seen as a minor trade-off for a greater upside – a solution to various modern-day datacenter power challenges that will not require additional PSUs, greater PSU capacity, or a loss in redundancy. As component density and quantity continues to increase with each generation, customers now require more and more power yet still have the same mechanical (limited space) or electrical (power budget) constraints. FTR resolves these challenges by allowing the total load to exceed the capacity of a single PSU during normal operation by utilizing the additional capacity of the redundant PSU, which results in a considerable increase in power standards and peaks during normal operating conditions.

That is what is so ironic about FTR – its "fatal flaw" of throttling has also become its "saving grace". FR does not allow for performance variations while FTR does, and this creates use cases where users can leverage FTR to support richer configurations without upgrading their PSU infrastructure. Figure 1 illustrates power, performance, and capability during normal operating conditions, while Figure 2 illustrates how power, performance, and capability during a PSU redundancy loss event:



Figure 1 – Example of FR/FTR performance during normal operating conditions

D&LLTechnologies



Figure 2 – Example of FR/FTR performance after PSU redundancy loss occurs

User Navigation Example

The latest-generation of PowerEdge servers (15G) support the option to choose Full Redundancy or Fault Tolerant Redundancy via PSU options at point of sale. Users can configure their servers via the sales portal on <u>www.dell.com</u> and have the option to click a step deeper via the Dell <u>Enterprise Infrastructure Planning</u> <u>Tool (EIPT)</u> for more granular guidance, as shown in Figure 3. Reviewing the PSU options in the PSU Guide and workload power details in EIPT will help PowerEdge users fine-tune their PSU configuration.

- **Gray** PSU capacity options cannot support the platform configuration
- White FTR. PSU capacity options can support the platform configuration, but peformance may be degraded after PSU redundancy loss
- **Green** FR. Minimum PSU capacity that can support the configuration with full performance during normal and fault operating conditions. Capacities greater than this capacity are also FR

Power Supply Guide					
					М
PSU Opt	ion a	Estimated Input A Power At Workload (W)	Estimated Maximum Potential Power (W)	Estimated Cost At An Workload (\$USD) (C Op Mi Po	nual Energy actional weration) at aximum Potential wer (\$USD)
	Dual, Hot-Plug, Fully Redundant Power Supply (1+1), 800W, Mixed Mode	767 W	1342 W	\$ 940.65	\$ 1645.83
•	Single, Hot-Plug, Power Supply Non Redundant (1+0), 800W, Mixed Mode	767 W	1342 W	\$ 940.65	\$ 1645.83
0	Dual, Hot-Plug, Power Supply Fault Tolerant Redundant (1+1), 800W, Mixed Mode	767 W	1342 W	\$ 940.65	\$ 1645.83
•	Single, Hot-Plug, Power Supply Non Redundant (1+0), 1100W, D, Mixed Mode	751 W	1350 W	\$ 921.03	\$ 1655.64
•	Dual, Hot-Plug, Fully Redundant Power Supply (1+1), 1100W, Mixed Mode	751 W	1350 W	\$ 921.03	\$ 1655.64
0	Dual, Hot-Plug, Power Supply Fault Tolerant 1100W Redundant (1+1), D, Mixed Mode	751 W	1350 W	\$ 921.03	\$ 1655.64
0	Dual, Hot-Plug, Fully Redundant Power Supply (1+1), 1400W, Mixed Mode	743 W	1321 W	\$ 911.22	\$ 1620.07
•	Rear Port Access, Single, Hot-Plug, Power Supply Non Redundant (1+0), 1400W, Mixed Mode	743 W	1321 W	\$ 911.22	\$ 1620.07
0	Dual, Hot-Plug, Power Supply Fault Tolerant Redundant (1+1), 1400W, Mixed Mode	743 W	1321 W	\$ 911.22	\$ 1620.07
0	Dual, Hot-Plug, Fully Redundant Power Supply (1+1), 2400W, Mixed Mode	751 W	1307 W	\$ 921.03	\$ 1602.9
0	Single, Hot-Piug, Power Supply Non Redundant (1+0), 2400W, Mixed Mode	751 W	1307 W	\$ 921.03	\$ 1602.9
0	Dual, Hot-Plug, Power Supply Fault Tolerant Redundant (1+1), 2400W, Mixed Mode	• 751 W	1307 W	\$ 921.03	\$ 1602.9
					Cancel

Figure 3 – Dell EIPT tool displaying various power and cost metrics based on configured PowerEdge server

D&LLTechnologies

For example, as seen in Figure 3, 2400W is required for FR. FTR enables the configuration to be supported with 1400W, 1100W, or 800W PSUs. If the platform were the R650 instead of the R750, the 2400W would not be an available option because it is the larger 86mm form factor which is not supported in the 1U 650. FTR enables this configuration to be supported when it could not be otherwise.

If the customer required the PSU input voltage to be low line AC (110V), the 1400W and 1100W PSUs would be limited to a 1050W output. The 2400W PSU would be limited to 1400W. Since 2400W is required for FR, this configuration could not be supported with FR. FTR enables this configuration to be supported with low line AC input.

EIPT estimates the typical power consumption with the 2400W PSU for the target workload to be 751W. The Maximum Potential Power (power virus) is estimated to be 1307W. Note, these are input power estimates, thus they are a little higher than the output power estimate and vary based on capacity due the PSU efficiency curves. The 2400W is the FR recommendation over the 1400W despite the worst case 1307W sustained power estimate because there are short duration power transients that exceed the 1400W power delivery capability.

FTR enables the customer to optimize CapEx and OpEx by right sizing their PSU capacity. 1400W could be an option to right size and still provide significant capacity to eliminate or minimize any potential performance degradation. With an estimated 751W typical power, the 1100W and 800W would be more aggress PSU right size options that provides the needed power for the user's workload assuming the workload does not change. If the workload or environment changes AND PSU redundancy is lost, FTR will manage the load increase to avoid unexpected shutdown and potential data loss.

Pros, Cons and Use Cases

Full Redundancy

- Pros
 - Consistent performance during normal operating conditions and PSU redundancy loss
 - No PSU throttling
- Cons
 - \circ $\,$ Maximum sustained power is constrained to the specifications of one PSU $\,$
 - Does not utilize the additional capacity of the redundant PSU during normal operation
- Use Cases
 - o Configurations that meet power requirements with only one PSU
 - \circ $\,$ Workloads that are sensitive to performance variations, such as HPC $\,$
 - o Platforms that do not have mechanical constraints, such as limited space for more PSUs
 - o Data centers that do not have electrical constraints, such as low-line AC

Fault Tolerant Redundancy

- Pros
 - o Allows for increased sustained perfomance during normal operating conditions
 - o Utilizes the additional capacity of the redundant PSU during normal operation
 - o Eliminates cost of purchasing additional or higher capacity PSU
 - Does not require giving up PSU redundancy
 - Does not require down-grading platform configuration to fit within target PSU capacity

D&LLTechnologies

- Cons
 - Performance may be reduced during PSU redundancy loss
- Use Cases
 - Configurations that would meet power requirements with the performance increase coming from the redundant PSU
 - o Platforms that have mechanical constraints, such as limited space for more PSUs
 - Data centers that have electrical constraints, such as low-line AC

Conclusion

Dell Technologies supports both Full Redundancy (FR) and Fault Tolerant Redundancy (FTR) options for the latest-generation (15G) of PowerEdge servers. By understanding the pros and cons of each redundancy type, users can optimize their server by upgrading or downgrading their configuration infrastructure based on what type of power redundancy they desire.



PowerEdge DfD Repository For more technical learning







Follow Us For PowerEdge news

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies and other trademarks are trademarks of Dell Inc. or its subsidiaries

D&LLTechnologies