

White Paper

PRIMERGY Modular Power Supply Unit

The Fujitsu modular power supply units enable reliable and efficient operation of the single and dual socket PRIMERGY servers. This white paper covers advantages, functions and technical details of the modular power supply concept.



Content

1 Modular power supply unit	2
1.1 Introduction	2
1.2 First and second generation modular power supplies	2
1.3 Fujitsu Battery Unit – An alternative to traditional UPS	2
2 Advantages of the modular PSU	3
2.1 Improved energy efficiency	3
2.2 Titanium and Platinum efficiency	3
2.3 Customer benefit of the 80 PLUS titanium PSU	4
3 Functions of the modular PSU	4
3.1 PSU operation at different input voltage levels	4
3.2 PowerSafeguard	4
3.3 Redundant modes	5
3.4 Power Management in iRMC / ServerView™	6
4 Selection criteria for the modular PSU	7
4.1 Focus on maximum performance and expandability	7
4.2 Focus on best efficiency	7
5 PowerCalculator in SystemArchitect:	8
5.1 Choosing the suitable power supply unit	8
5.2 Measured and calculated power consumption vs. power rating plate	8

1 Modular power supply unit

1.1 Introduction

The modular power supply unit (abbreviation: modular PSU) is a common power supply concept for the vast majority of Fujitsu PRIMERGY tower and rack servers. The space-saving 1U power supply unit (185 mm x 73.5 mm x 40 mm) is available in different versions. With the launch of the dual M1 PRIMERGY server family in autumn 2014 Fujitsu introduced a second generation of power supply units which support the "turbo mode" of the latest systems. It is not possible to mix power supplies between 1st and 2nd generation as there is no up- and downward compatibility. Also mixing of different types of one PSU generation is not supported.

- 450 W, 80 PLUS platinum level, 1st and 2nd generation
- 800 W 80 PLUS platinum level, 1st and 2nd generation
- 800 W 80 PLUS titanium level, 2nd generation
- 1200 W 80 PLUS platinum level, 2nd generation
- Fujitsu Battery Unit 280 W, 1st generation
- Fujitsu Battery Unit 380 W, 2nd generation



The modular PSU is available for a range of PRIMERGY servers:

PRIMERGY	450 W Platinum (94 %)		800 W Platinum (94 %)		800 W Titanium (96 %)	1200 W Platinum (94 %)	Fujitsu Battery Unit	
	1st gen	2nd gen	1st gen	2nd gen	2nd gen	2nd gen	1st gen	2nd gen
RX1330 M1	✓						✓	
RX1330 M2		✓						✓
RX2520 M1	✓		✓		✓			
RX2530 M1		✓		✓	✓			
RX2540 M1		✓		✓	✓	✓		
RX2560 M1		✓		✓	✓	✓		
TX1330 M1	✓						✓	
TX1330 M2		✓						✓
TX2540 M1	✓		✓		✓			
TX2560 M1		✓		✓	✓	✓		

PRIMERGY dual socket servers generally support both the 450 W power supply unit and the 800 W power supply unit. In contrast, the mono socket servers have lower power consumption and therefore only require the 450 W power supply unit. The 1200 W power supply is currently only supported by the PRIMERGY RX2540 M1 and TX/RX2560 M1.

The modular PSU also supports the proven security function called "PowerSafeguard", which even safeguards further operation of the server if a power supply unit fails during redundant operation and the power consumption exceeds the capacity of the remaining power supply unit. Depending on the server configuration, this protection is available up to an excess capacity of 30 %. (See chapter 3.2)

1.2 First and second generation modular power supplies

To support the latest Intel® Turbo Boost Technology in the Dual M1 PRIMERGY servers the second generation of modular power supply units was introduced in 2014. This new PSUs will become the standard for the upcoming single socket PRIMERGY M2 servers. Mixing first and second generation power supplies is not supported.

1.3 Fujitsu Battery Unit – An alternative to traditional UPS

The Fujitsu FJBU internal battery backup unit is a cost-effective, reliable and easy solution to secure data in a power outage. The compact device fits into the modular PSU slot and – in case of a blackout - keeps the server running to achieve at least a graceful shutdown. Its Ni-MH battery allows for a very long life time of up to 5 years. Currently it is released for the PRIMERGY TX1330 M1 / M2 and PRIMERGY RX1330 M1 / M2.



2 Advantages of the modular PSU

The modular PSU offers a range of advantages, in particular in the areas of energy efficiency and easier maintenance and stocking.

2.1 Improved energy efficiency

High energy efficiency depends on the utilization level of the power supply units. The best possible degree of efficiency is only achieved with a utilization level of 50 %. Above 50 % the efficiency slightly decreases again (see table below). The server configurations are often very different. The energy consumption spectrum ranges from 60 Watt up to more than 1000 Watt. To obtain the best possible energy efficiency with every configuration, the modular PSU concept offers units in different sizes so as to enable the individual configuration of the PRIMERGY servers with an adequate power supply. The modular PSU concept also enables an individual configuration for phase and PSU redundancy with an optimum degree of efficiency of the power supply units.

Modular power supply units rely on technology with the highest energy efficiency standard, which enable an efficiency degree of up to 96 % and thus meet the CSCI certification standard "80 PLUS titanium". This high degree of efficiency ensures that merely 4 % of the transformed energy is lost in the form of waste heat and is not available to the server.

Furthermore, PRIMERGY Power Management - with its easy and automated energy management - helps achieve significant savings in energy requirements. For example, innovative functions, such as the automatic deactivation of unused PCIe ports, help to further reduce power requirements.

The combination of efficient, modular power supply units with intelligent power management contributes to making the PRIMERGY servers among the most efficient in their class. This for example has enabled the Fujitsu PRIMERGY RX2560 M1 to be ranked as the second most energy efficient server according to the SPECpower benchmark.¹



2.2 Titanium and Platinum efficiency

Modular 450 W, 800 W, and 1200 W power supply units fulfill the state-of-the-art '80 PLUS platinum' efficiency class. With the 800 W 80 PLUS titanium power supply units Fujitsu offers the industry leading highest efficiency. The degree of efficiency depends on the utilization level of the power supply unit. The best possible degree of efficiency is therefore only achieved with an approx. 50 % utilization level, as illustrated in the table below. Unnecessary oversizing of power supply units thus reduces the possible energy efficiency.

80 PLUS Levels	10 % Load		20 % Load		50 % Load		100 % Load	
	Efficiency	PF	Efficiency	PF	Efficiency	PF	Efficiency	PF
Silver	75	0.65	85	0.8	89	0.9	85	0.95
Gold	80	0.65	88	0.8	92	0.9	88	0.95
Platinum	82	0.65	90	0.8	94	0.9	91	0.95
Titanium	90	0.65	94	0.95	96	0.9	91	0.95

(230 V Internal Redundant, PF = Power Factor)

¹ Competitive benchmark results stated above reflect results published on www.spec.org as of February 12, 2016. Score of 10.699 overall ssj_ops/watt published by Fujitsu on the PRIMERGY RX2560 M1, see: http://www.spec.org/power_ssj2008/results/res2015q2/power_ssj2008-20150317-00690.html. SPEC and SPECpower are trademarks or registered trademarks of Standard Performance Evaluation Corporation (SPEC). For the latest SPECpower_ssj2008 results, visit http://www.spec.org/power_ssj2008/results/power_ssj2008.html.

2.3 Customer benefit of the 80 PLUS titanium PSU

While the 80 PLUS titanium PSU costs a little more than the platinum PSU, this additional costs will be over-compensated within the regular lifetime of a server. Like the calculation below shows, over a lifetime of 5 years a titanium PSU can save up to 275 € in energy costs.

Example

- PRIMERGY RX2540 M1, redundant 800 W PSU, PSU load 25 % equals 200 W net power consumption per PSU, 400W overall power consumption
- 80 PLUS Titanium PSU: 12,7 W loss (94 % efficiency @ 25 % load)
- 80 PLUS Platinum PSU: 22,2 W loss (90 % efficiency @ 25 % load)
- Average power saving: 19 W for 2x PSU
- Server runtime: 24 hours / 365 days
- Productive lifetime: 5 years
- Energy costs: 0,15 € / kWh
- Power usage effectiveness (PUE): 2.2

⇒ $19W * 24h * 365d * 5a * 0,15€ * 2,2(PUE) = 275 €$

Possible energy savings: 275 € for 5 years

The 80 PLUS titanium power supply unit underlines Fujitsu’s overall strategy to provide energy efficient solutions for IT infrastructures.

3 Functions of the modular PSU

3.1 PSU operation at different input voltage levels


While the “bread and butter” power supply units with 450 W / 800 W platinum can be operated at a very wide range input voltage, the 800 W titanium and 1200 W PSUs have some limitations. The 800 W titanium modular power supply with 96 % energy efficiency has to be operated in high line voltage range. For the 1200 W PSU, the maximum output wattage is reduced under low line voltage conditions. Find details and nominal voltage levels in the table below.

PSU Types	100 - 109 V	110 - 130 V	200 - 240 V
450 W platinum (1 st , 2 nd gen)	450 W	450 W	450 W
800 W platinum (1 st , 2 nd gen)	800 W	800 W	800 W
800 W titanium	-	-	800 W
1200 W platinum	900 W	1000 W	1200 W

3.2 PowerSafeguard

The PowerSafeguard protects continuous operation of a server if a power supply unit in a non-redundant 2x PSU configuration fails. In such a case the PowerSafeguard is able to throttle the system performance and energy consumption within a few milliseconds to keep the system working with just a single power supply. The PowerSafeguard guarantees an uninterrupted operation of server configurations with max. 30 % higher energy consumption than one single power supply can provide.

To do this PowerSafeguard sets the working frequency of the processor to a minimum and then successively increases it again until the maximum performance of the power supply unit is reached. However, the resulting impact on the entire performance in a high configuration is limited in most cases, as only short-term peak utilization is no longer supported, as the diagram above shows. Most common average utilization of servers is at approx. 30 %, in virtualization environments approx. 50 - 70 %.



Example:

- Dual socket server with 550 W energy consumption
- 2x 450 W modular power supply units

Without PowerSafeguard

As the PSUs in this configuration are not redundant, the server would stop working if one of the PSUs fails. The result: data loss, failure of important software, serious performance problems with the remaining servers, etc.

With PowerSafeguard

Even in this non-redundant configuration the PowerSafeguard will keep the server working if one of the PSUs fails. PowerSafeguard throttles instantly the system performance. As a result the energy consumption will decrease from 550 W to the 450 W maximum the single remaining PSU can provide. The server will keep on working, no data loss has to be feared and software systems will continue to operate with slightly less performance at worst.

The PowerSafeguard is available in each PRIMERGY server with support for 450 W, 800 W, and 1200 W modular power supplies.

3.3 Redundant modes

Modular PSUs support a range of redundancy modes:

- **Non-Redundant Mode**
In this mode, the PSU output is sufficient to support the installed system configuration. If one power supply unit fails, the system may shutdown immediately.
Possible configurations: 1+0
- **PSU-Redundant Mode**
In this mode one spare PSU is available, so if one PSU fails, the system continues operation.
Possible configurations: 1+1

If one PSU fails, peak system performance can be reduced by PowerSafeguard, if a smaller PSU was chosen to enable more efficient operation. For applications where full system performance even during the exchange of the faulty PSU is needed, the larger 800W or 1200W PSU should be selected.

- **AC Redundant (Dual Phase) Mode**
In this mode the power supplies will be connected to two different power lines, so that the system can continue to work even if a power line or a single PSU fails.
Possible configurations: 1+1

Redundancy modes of PRIMERGY servers with 2 power supply units:

Config	PSU 1	PSU 2	Up to power consumption	Redundancy level
1	450 W	-	450 W	1+0
2	450 W	450 W	450 W	1+1
3	450 W	450 W	802 W	2+0
4	800 W	-	800 W	1+0
5	800 W	800 W	800 W	1+1
6	800 W	800 W	1428 W	2+0
7	1200 W	-	1200 W	1+0
8	1200 W	1200 W	1200 W	1+1
9	1200 W	1200 W	2148 W	2+0



A mixed mode between 450 W, 800 W Platinum, 800 W Titanium and 1200 W is not released for continuous operation. However a mixed operation can be used for upgrades, servicing and short periods of time.

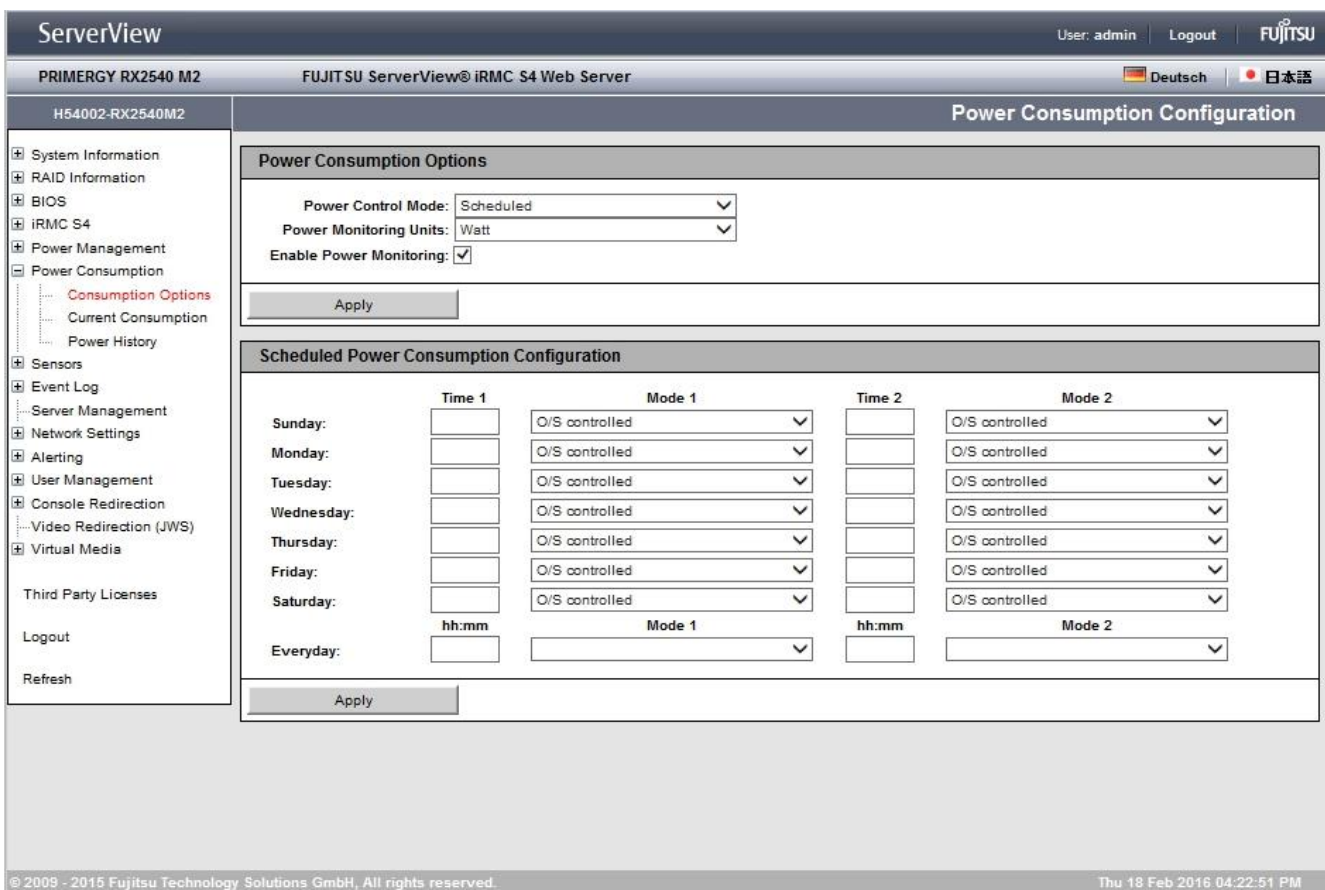
3.4 Power Management in iRMC / ServerView™

Power Management in iRMC/ ServerView™ permits with its various energy profiles the energy requirements of the PRIMERGY server to be adapted according to individual requirements.

Several pre-defined energy modes are available for this purpose:

- O/S controlled:
The iRMC S4 leaves the control of the energy consumption to the operating system.
- Minimum Power:
The iRMC S4 controls the server with regard keeping energy consumption as low as possible. In this case, performance is not always at an optimum. This mode is especially interesting for times of low performance requirements (weekends/nights), thus enabling significant savings to be made in energy requirements.
- Power Limit:
A pre-defined energy consumption will not be exceeded. This mode is required, if specific energy values in the rack or in the data center compulsory must not be exceeded.
- Scheduled:
The individual energy modes can also be assigned to a time schedule, which determines the time and day on which an energy mode is activated (see figure below).

The individual energy modes can also be assigned to a time schedule, which determines the time and day on which an energy mode is activated (see screenshot below).



Please note: Above settings are CPU specific and are not available for all CPU types and Fujitsu PRIMERGY servers

More information about power management and its options is available here:

- German: <http://manuals.ts.fujitsu.com/file/11472/irmc-s4-ug-de.pdf>
- English: <http://manuals.ts.fujitsu.com/file/11470/irmc-s4-ug-en.pdf>

4 Selection criteria for the modular PSU

When selecting a suitable power supply unit, you should first determine on which value focus should be placed. Two possible options are:

- Maximum performance and expandability
- Optimal efficiency

4.1 Focus on maximum performance and expandability

If emphasis is placed on maximum performance (at all times) and/or on the maximum expandability of the server for future requirements, then the PRIMERGY dual socket server should be equipped with the 800 W or 1200 W modular PSU. This ensures full system performance with each configuration, even if a power supply unit in a 1+1 configuration fails.

PowerSafeguard does not intervene here.

4.2 Focus on best efficiency

If focus is placed on the achievement of best possible energy efficiency in the data center, then two essential aspects should be taken into account when selecting the power supply unit:

- Configuration of the PRIMERGY server and
- Future upgrades

Remember: The degree of efficiency depends on the utilization level of the power supply unit. The best possible degree of efficiency is therefore only achieved with an approx. 50 % utilization level (see chapter "2.2 Titanium and platinum efficiency")

The specific configuration of the PRIMERGY server is decisive for the selection of the suitable power supply unit. The following components thus have a considerable influence on overall power consumption:

- Number and class of Intel processors
- Number and type of hard disks (2.5"/3.5"; SSD/SATA/SAS; 7.2K/10K/15K)
- Number and type of connected I/O cards
- Number and type of connected RAM modules
- Configuration of GPGPU, graphic, or computing cards

The following overview mirrors the energy consumption of a RX2540 M1 with 100 % CPU utilization in the respective configurations.

PRIMERGY RX2540 M1: Estimated Power Consumption <small>(For illustration only – use PowerCalculator to check your PRIMERGY configuration)</small>		Memory configuration							
		1 x 8 GB	4 x 8 GB	1 x 8 GB	4 x 8 GB	4 x 16 GB	12 x 16 GB	12 x 16 GB	
		2 x 8 GB	8 x 8 GB	2 x 8 GB	8 x 8 GB	8 x 16 GB	24 x 16 GB	24 x 16 GB	
CPU configuration	1 CPU	1 x E5-2630L v3 55 W	163 W	176 W	216 W	229 W	368 W	411 W	740 W
		1 x E5-2630 v3 85 W	187 W	199 W	243 W	255 W	395 W	438 W	760 W
		1 x E5-2670 v3 120 W	231 W	244 W	298 W	311 W	451 W	494 W	n/a
		1 x E5-2690 v3 135 W	252 W	264 W	320 W	333 W	474 W	517 W	n/a
	2 CPU	2 x E5-2630L v3 55 W	231 W	257 W	285 W	310 W	455 W	543 W	867 W
		2 x E5-2630 v3 85 W	279 W	304 W	337 W	363 W	508 W	600 W	914 W
		2 x E5-2670 v3 120 W	369 W	395 W	450 W	476 W	629 W	718 W	n/a
		2 x E5-2690 v3 135 W	411 W	437 W	495 W	521 W	676 W	766 W	n/a
		2 x 3.5", 2 x PCIe		8 x 2.5", 4 x PCIe		24 x 2.5", 6 x PCIe		16 x 2.5", 2 x Nvidia Grid, 4x PCIe, Blu-ray, LTO6	
		HDD and PCIe configuration							

Recommendation for best energy efficiency & redundancy

- 2 x 450 W PSU recommended – no impact if 1 PSU fails
- 2 x 450 W PSU possible – system performance will be throttled down by PowerSafeguard if 1 PSU fails
- 2 x 800 W PSU is recommended
- 2 x 800 W PSU possible – system performance will be throttled down by PowerSafeguard if 1 PSU fails

Even if there is no simple formula for selecting a 450 W or 800 W power supply unit, it can however be said that if the following sample configuration is met or exceeded an 800 W power supply should be chosen:

Sample configuration
<ul style="list-style-type: none"> ■ 2x Intel Xeon E5-2670 v3 ■ 8x Memory DIMMs ■ 24x HDDs ■ 6x PCIe cards



As a matter of principle, the respective **PRIMERGY configuration should be configured in SystemArchitect and the power consumption checked using the PRIMERGY PowerCalculator** (see section 5) in order to obtain a sound basis for selecting a suitable power supply unit. And in the event of a pending server upgrade, the configuration should first be checked in the PowerCalculator to determine whether a larger power supply unit is required.

5 PowerCalculator in SystemArchitect:

5.1 Choosing the suitable power supply unit

The PowerCalculator in SystemArchitect provides excellent help when it comes to calculating the performance values and thus with the selection of the suitable power supply unit.

Symbol	Status	Description	Load [%]	DC-load [W]	Active Power [W]	Apparent Power [VA]	Current [A]	Heat Emission [kJ/h]	Air Flow Rate [m3/h]	Price / kWh	Costs / Year
		PY RX2540 M1 24x 2,5"	100	705,40	766,47	769,54	3,34	2.759,29	96,00	0,15 €	1.007,14 €
					766,47	769,54	3,34	2.759,29	96,00		1.007,14 €

A detailed description of how to calculate the performance values and which points should be taken into account are available under: English, German: <http://configurator.ts.fujitsu.com/public/information>

5.2 Measured and calculated power consumption vs. power rating plate

The calculated power consumption of the PowerCalculator is derived from a database, accumulating single power values of all components of a configuration. The workload profile for the measured database values is the same as the SPECpower benchmark profile, and the CPU load can be configured from idle till 100 %.

In reality system workloads can be different and most servers operate by far below 100 % CPU load, which leads to a lower power consumption as the calculated power consumption.

The value on the power rating plate is derived to fulfill regulations and for worst case sizing of electrical facilities and is higher than calculated or typical measured power consumption. If accurate power consumption is needed for projects, tenders etc. – please contact primergy-pm@ts.fujitsu.com.

Contact

FUJITSU CEMEA & I
 Address: Bürgermeister-Ulrich-Str. 100, 86199 Augsburg,
 Germany
 E-mail: florian.frimmel@ts.fujitsu.com
 Website: www.fujitsu.com/global
 2016-02-18 EN

© 2016 FUJITSU CEMEA & I. Fujitsu, the Fujitsu logo, and ServerView are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.