

WHITE PAPER

Cool-safe™: what is it and why is it good for you?

Issue April 2009

Pages 7

Contents

| | |
|---|----------|
| Introduction | 1 |
| Performance – what your CEO loves best | 2 |
| Efficiency – why your CFO loves PRIMERGY servers | 2 |
| Reliability – don't make your CFO halve your lifetime! | 4 |
| Conclusion | 7 |

Introduction

Cool-safe is a holistic system engineering methodology used in PRIMERGY servers. Cool-safe is right at the heart of our engineering efforts. If you ask a group of our engineers to say what Cool-safe means to them, you would come up with these 3 common denominators:

- Performance
- Efficiency
- Reliability

Our PRIMERGY engineers are putting all their creativity at your service in order to ensure that efficiency, reliability and performance runs right through your PRIMERGY server.

So, that's what Cool-safe is in a nutshell. Oh yeah, Cool-safe is also about keeping your CEO and CFO happy. Now's let's take a look inside and see how this is done in practice...the engineering.

Performance – what your CEO loves best

Your CEO wants to be able to take advantage of every business opportunity out there and PRIMERGY engineers know that. That is why they use the Cool-safe method to get the best performance out of the system, without losing sight of the efficiency and reliability. A particularly important development in this respect is the introduction of the new Intel® Nehalem processor generation with high performance features such as Intel® QuickPath Interconnect. On top of that PRIMERGY S5 systems are capable of using up to 144Gb DDR3 memory modules with 1333MHz, and up to 20 hard disks! Now let us look at how to keep this beast under control.

Efficiency – why your CFO loves PRIMERGY servers

Efficiency is a big word and it can mean many things, but in this context efficiency means not wasting energy. Wasting money is never a good idea, and this is certainly the case these days when the world economy is in bad shape. Wasting energy is also bad from an environmental point of view too. Both of these are very important when it comes to data centers where you have many servers running in rack or blade infrastructures. Did you know, for example, that the consumption of fully populated blade system can reach a staggering 3600 Watts?

Now let's take a look at what makes PRIMERGY servers efficient.

Power supply unit (PSU) efficiency

In our S5 systems we build in PSUs capable of reaching power efficiency over the 80% mark. That means that only some 20% of the energy going through the power supply is lost. Now, if you were a technical whizz kids, you might ask 'but at what load range is that then?' I would be very impressed by your know-how, and yes, you're right. Energy efficiency doesn't mean much unless you relate it to the workload going through the PSU. Rest assured, these amazing power efficiency levels can be reached with workloads ranging from as little as 15% to as much as 85%.

BIOS and Processors

How does the BIOS affect energy consumption? Well there's no point in the hardware running at full speed all the time. Your local energy supplier doesn't mind, but your CFOs does. That's why processors have so called "C-states" and "P-states". C-states are idle states and P-states are operational states. In other words, the c-state decides how for a processor should go into idle mode. The deeper the idle state, the less energy it consumes but the longer it can take to 'wake up'. The p-state decides at what frequency the processor should be run. By running at a lower speed, the processor generates less heat and consumes less power. In Intel® processors, P-states are known as Enhanced Intel Speed Step® Technology. Still with me? Of course, the processor interacts with the operating system in order to set the right c or p state, but that's not the whole story. Here's where the BIOS engineers come in... The BIOS has a table of values which sets the overall limits within which the hardware and software operate. That's what determines the how low your system is idling or how high it is performing. Not an easy task, but that's where the PRIMERGY engineering crew can blend in their awesome analytical skills with their fearsome precision. The BIOS know-how is the result of extensive testing and customer feedback.

Motherboard design

So, what have these PRIMERGY engineers thought about here then?

Well, to achieve a high efficient server system you have to look at the power delivery concept as a whole, from AC Power in down to the guys on the motherboard, like processor, memory and chipset, using this energy to do their work.

In our brand new RX/TX 300 S5 systems, they have opted for a complete redesign of the whole power delivery structure. Always keeping energy efficiency as a priority one goal in mind, they first replaced the old power supplies by new ones, optimized for high energy efficiency.

Further, they have also been able to integrate the power backplane completely onto the motherboard. The power backplane was formerly used to convert the 12V input power, coming from the redundant power supplies, into the main voltages needed usually by a server motherboard (3,3V, 5V, 5Vstby, -12V). Additionally the power backplane contained the current limiting circuits for each of the six different 12V rails. The connection from the power backplane to the motherboard was then achieved by over 30 power cables attached by large connectors on each end. As high currents over connectors and cables produce energy losses, this power backplane solution contributed significantly to lower energy efficiency. The new, highly efficient redundant power supplies are now connected directly to the motherboard via low impedance connectors. This guarantees a highly efficient energy transport from the power supply to the motherboard.

Last but not least, the PRIMERGY crew came up with a completely redesigned onboard power delivery concept.

Now, the power converters are arranged in parallel and the loss loaded cascading of the regulators is now eliminated. Additionally the use of highly efficient DCDC switching converters replaces old fashioned linear converters.

All those changes lead to a very high energy efficient server system. Besides that, the power cabling is reduced which improves the serviceability as well. Have a look. Here's what the insides of an RX300 S5 looks like when you take the cover off.



PRIMERGY ServerView Management Suite

All PRIMERGY servers are delivered with the ServerView Management Suite. The ServerView Management suite is a tool to setup, manage, integrate and recover IT infrastructures. We want to look at the management part. The secret lies in the iRMC2 chip (integrated Remote Management Controller) Integrated on the motherboard (see below). Now let's see what the iRMC2 chip and ServerView can do together.



Power Control

This money-saving functionality is one your CFO's favourites. The integrated Power Management controls (remote) power on, power off and shutdown for physical servers as well as on virtual servers.

| View Name | Query Name |
|--|-------------------|
| SV_PowerConsumptionServerPerDay | PowerCons1 |
| SV_PowerConsumptionServerPerLast30Days | PowerCons30 |
| SV_PowerConsumptionServerPerYear | PowerCons366 |
| SV_PowerConsumptionServerPerLast7Days | PowerCons7 |
| SV_PowerConsumptionServerPerLast90Days | PowerCons90 |
| SV_PowerConsumptionServerOverall | PowerConsAll |
| SV_PowerConsumptionServerLastHalfYear | PowerConsLHYear |
| SV_PowerConsumptionServerLastMonth | PowerConsLMonth |
| SV_PowerConsumptionServerLastQuarter | PowerConsLQuarter |
| SV_PowerConsumptionServerLastWeek | PowerConsLWeek |
| SV_PowerConsumptionServerLastYear | PowerConsLYear |

Power Monitoring

Another CFO favourite. The Power Monitor enables the analysis of the power consumption of a single or multiple servers. The consumption data are logged for settable periods, and displayed as graphs.



Power Consumption Control

Thanks to the iRMC2 chip, you have a choice between two modes:

- Minimal Power Consumption: you constrain the CPU's power to lowest minimum used by the CPU type.
- Best Performance: all of the CPU's performance is available for the operating system.

The scheduler function lets you switch between these modes on a planned basis. Last but not least, should the uninterruptible power supply be activated because of a power cut the server reduces the power consumption to a minimum.

Power Consumption Tool – fine tuning the consumption

The end user also has the chance to fine tune their consumption by using the "Power Consumption Tool". The tool is available from the [Extranet](#). It simulates load factors for all server configurations and displays the results immediately. Here is an example based on the PRIMERGY TX150 S6. As you can see, you can select various parameters and the tool shows you the power consumption for the chosen configuration. You can also use this great tool to help you plan the layout of a data center. The details here come from extensive testing procedures and years of experience in delivering efficient systems. The engineers have tested many different setups so that you can choose the one that minimizes consumption.

That's why your CFO loves PRIMERGY engineers. Because they do everything possible to minimize waste and save your company the very last cent. And because it makes him feel good that he can help preserve the environment too.

| TX150 S6 | |
|-------------------|---------------------|
| CPU Type | Xe E3110 UP 3.00GHz |
| CPU Quantity | 1 |
| Memory Type | DDR2 - 800 |
| Mem Capacity | 1x2 GB |
| HDD Type | SAS (2.5") |
| HDD Quantity | 2 |
| Accessible Drives | (--) |
| PCI Card Quantity | (--) |

| Summary TX150 S6 - Configuration 1 | | |
|------------------------------------|----------|------|
| Measures | Metric | US |
| Voltage | 230 | V |
| Ambient Temperature | 10 C-25C | |
| Liquid Cooling Rack | No | |
| Weight | 21,2 | Kg |
| Apparent Power | 71 | VA |
| Active Power | 63 | W |
| Current | 0,3 | A |
| Number of Cords | 1 | |
| Max Current per Cord | 0,3 | A |
| Heat Emission | 226 | kJ/h |
| Air Flow Rate | 40 | m3/h |

Reliability – don't make your CFO halve your lifetime!

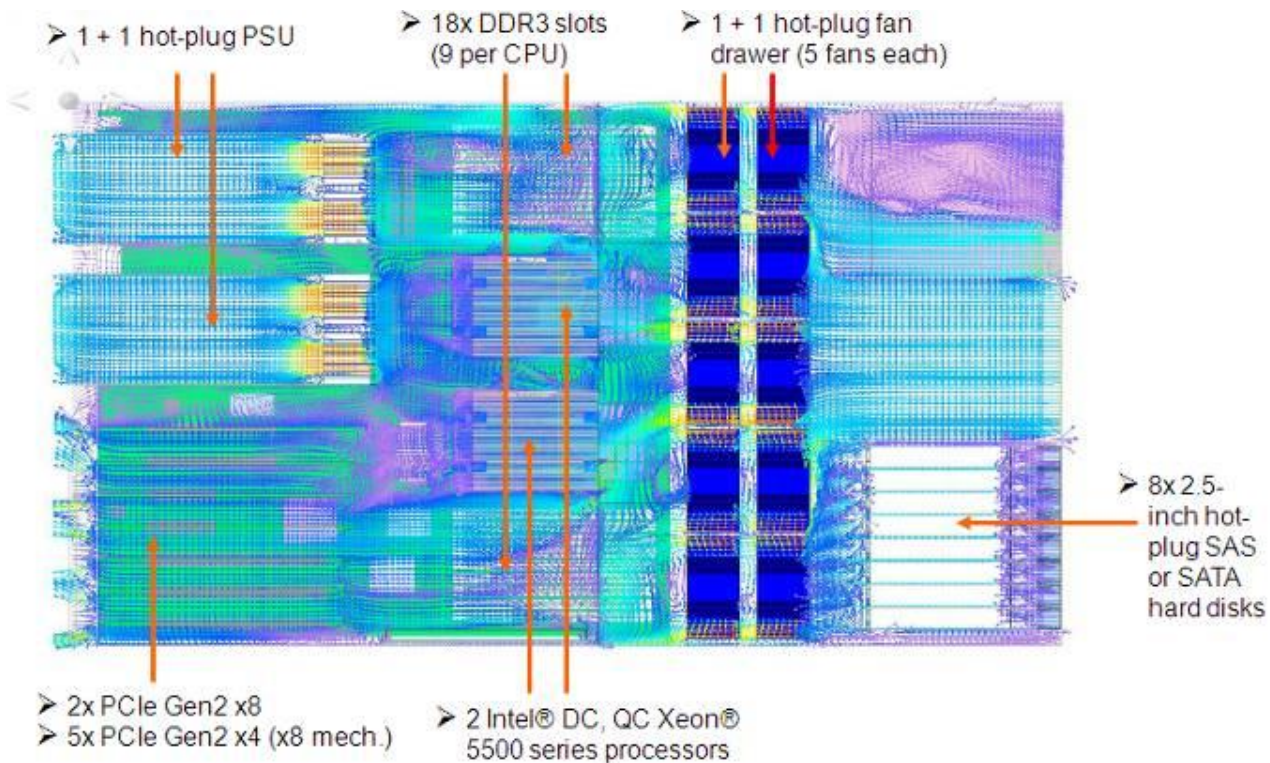
An unreliable system costs time and money. That is why PRIMERGY engineers do everything possible to make sure that the server keeps on going. Before we go on, keep 3 key points in mind:

1. 10°C more heat halves the MTBF (Mean Time Between Failure)
2. Reducing temperature of electronic components by 15°C doubles component lifetime
3. Reducing temperature of mechanical components by 10°C doubles component lifetimes

This suggests that there is a clear correlation between operating temperatures and failure rates. Or try looking at it this way, if your system keeps failing, your CFO may think about halving your lifetime!

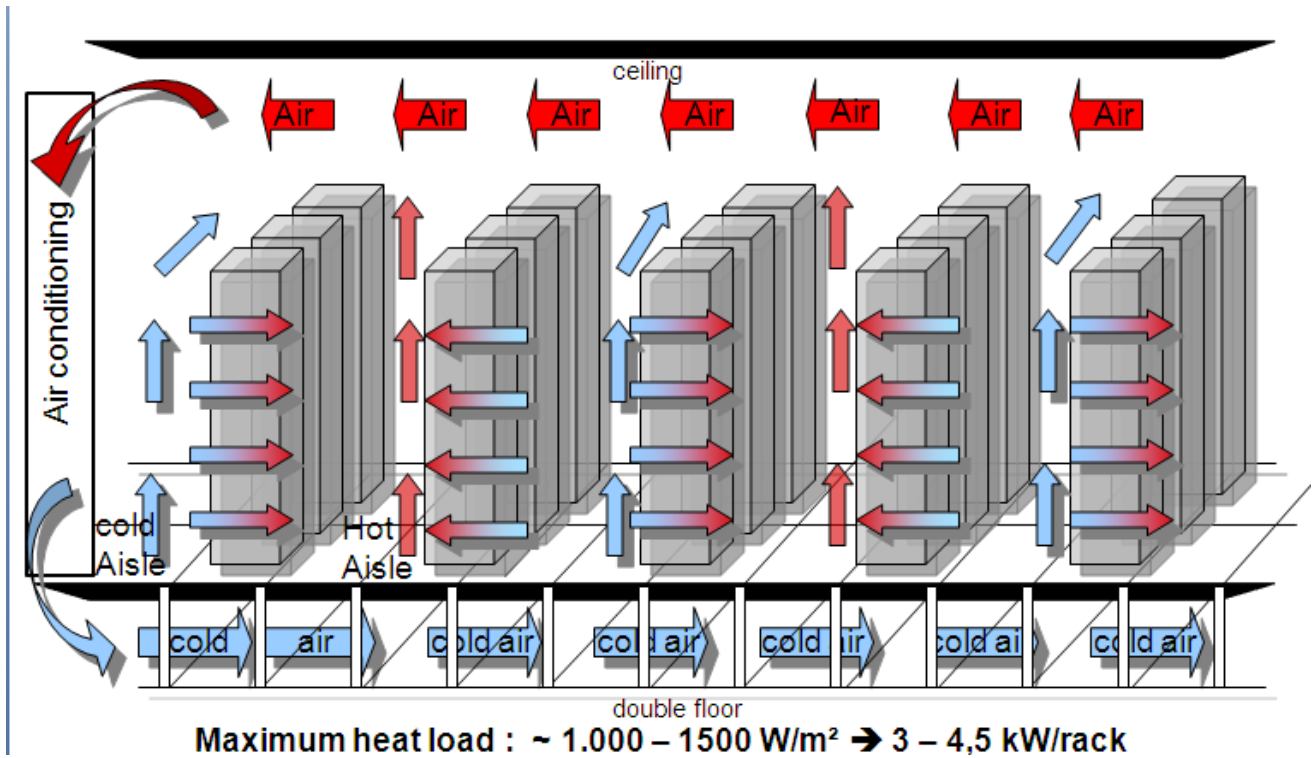
Air flow – keeping the air cool

This is absolutely key when it comes to reliability. The objective here stop the system from overheating by keep cool air coming into the system, and moving the hot air out of the system. In order to simulate the air flow, PRIMERGY engineers use Computational Fluid Dynamics (CFD). CFD is widely used in such high-tech fields as aeronautics. The picture below gives an example of a CFD applied to a PRIMERGY RX300 S5. The blue areas are the cooler areas and the red ones the warmer ones.



The air flows through the system from right to left. As it comes in, the air is channeled into the system via the fans distributed along the length of the chassis. The most critical components are strategically placed directly behind the fans, guaranteeing optimum cooling. The power supplies have their own fans (look at the yellow rectangles at the top left of the picture), which is not really surprising. After all, by generating the power necessary to keep your system up and running, and this generates heat which has to be removed from the system.

This effort is necessary on the one hand in order to cool a free standing system, such as a tower, and also when the system is in a rack in a data center. The picture below shows the air flow in a data center.




As you can see from the picture, the systems have to be engineered in such a way that:

- A single rack or blade infrastructure is cooled appropriately and
- that it does not affect the reliability of other racks or blades

The other key point from the drawing is that the large amounts of heat, mean a lot of effort spent in keeping the data center cool. This of course relates closely to the topic of efficiency and saving money.


So, on the one hand the PRIMERGY engineers look at the system as a single unit and also as part of a bigger picture. In order to do that, the engineers have to keep one basic principle in mind: a clean design. And this is how they do it... The secrets of a cool system.

The picture below of a rack system gives an idea of how it is setup in order to guarantee a cool system.



1. There is an unobstructed airflow, with large enough 'corridors' for the air to pass through. Notice also that the engineers have done their best in order to keep the design nearly cable less. This minimizes any possible obstructions.
2. Heat is quickly moved away from hot spots
3. The air flow is conducted to where it is needed most. Look at how the fans are setup across the width of the system. The cooling channels are kept separate in order to prevent shortcuts and loops in the airflow

Another great example is given by the TX300 S5. The air speed is kept as low as possible by using large fans turning slowly.




In this case the 3 redundant fans are blowing the air directly towards the processors and memory slots. The large housing and the big fans reduce the system temperatures. The large housing allows more air to circulate through the system, making the system cooler. Including many small fans, would mean that they would have to turn very fast in order to help keep the system cool. These small fans would make the system power-hungry and noisy. That is why the TX300 S5 uses 3 big fans instead. These turn slower, making the system more energy efficient and more silent.



A large diameter fan produces good air flow even at low number of revolutions per minute. On top, the PRIMERGY servers feature the ServerView Management Suite, which also controls the fan. The speed is adjusted to the needs, between

- Full speed specification: 3500 rpm (+/- 10%)
- Low speed specification: 1400 rpm (+/- 10%)

The fan speed is controlled in relation to the ambient temperature and the server power consumption. This can be viewed in graphical form in ServerView.



Of particular importance for tower systems is the new honeycomb design of the front bezel. While this may not seem like a great innovation at first sight, the hexagonal holes let in more cool air than circular ones. This increases the system's reliability. It also increases the performance as it allows the system to run at higher speeds. It's a small but effective improvement.

Conclusion

So, now the engineering lesson is over and now you know what makes up the Cool-safe methodology, what do you make of all this?

Performance: Cool-safe lets you take full advantage of the performance at your fingertips. A system which is already overheating at a low performance load, just won't be able to perform any more. It has hit a brick wall. A system which is cooled appropriately, will be able to perform at higher performance loads.

Efficiency and Green IT: Cool-safe saves money and contributes to the environment. By keeping the system running at a low temperature range you avoid inefficient use of resources, both financial and natural.

Reliability: Cool-safe saves headaches. By keeping the system cool and preventing it from over-heating it preserves the life of your components. One of the main reasons why components fail is simply that they overheat. Cool-safe works against this, which means that your system lifetime is enhanced.

Pure innovation: that is the engineering lifeblood running through Cool-safe.