Moving Towards 5 Kilowatt Boards

The twin server-rack challenges of supplying power and removing resultant heat means that both evolutionary advances and dramatically new approaches may be viable; only time will tell which proves to be the winning approach. **Patrick Le Fèvre, Marketing Director, Ericsson Power Modules, Sweden**

Power demand per board in the data-server segment has increased by a factor of four since the early 1980s, from just 300 W to 1200 W. It is further forecast to grow to 3 kW per board by 2015, with some studies forecasting as high as 5 kW/board consumption by 2020. This relentless trend brings two major challenges to power designers: how to develop highly efficient topologies and products to deliver this level of power from isolated bricks, and how to cool the circuit boards efficiently to ensure reliable, long-life operation.

According to sources such the Ericsson Mobility Report (www.ericsson.com/mobility-report) the amount of data transferred through networks is increasing at a tremendous rate. The mobile industry generates about 1.9 exabytes/month and it is estimated this will increase by a factor of ten between 2013 and 2019. Data centers are seeing corresponding growth: it is estimated that annual IP traffic was 2.6 zettabytes in 2012, and will be at 7.7 zettabytes by the end of 2017.

To respond to this demand and its power needs, data-center operators are considering two approaches such as modernization of existing data farms or new data centers using the latest technologies for higher energy utilization and efficiency, resulting in a reduction in both total cost of ownership (TCO) and environmental impact.

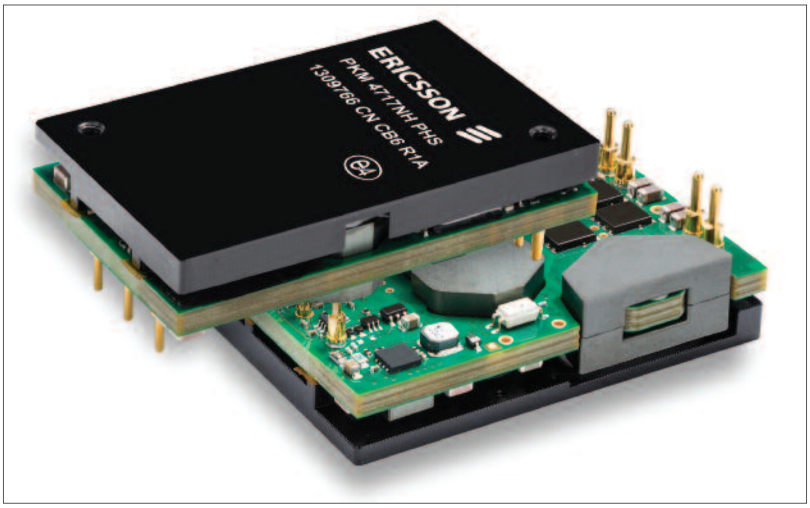

Figure 2: **Datacom-input-range high-power-density 864 W quarter-brick module featuring double power pins and optimized design for cold-wall cooling**

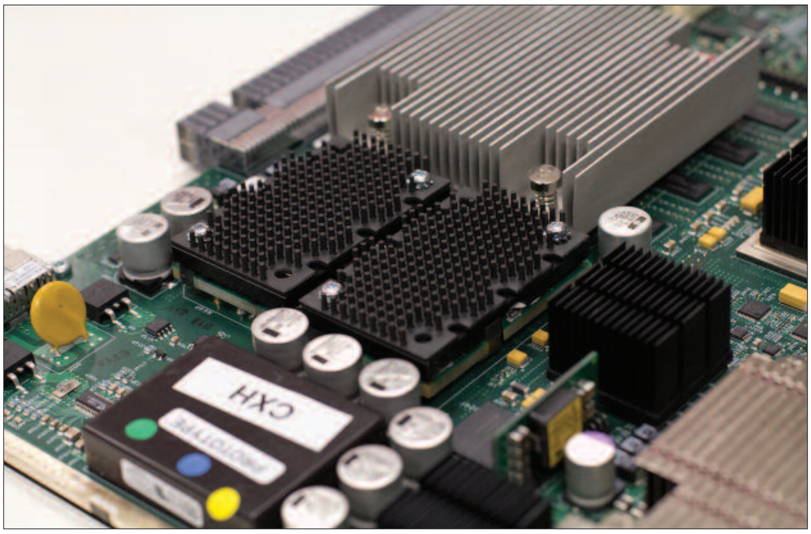

Figure 1: **Today, two quarter-bricks in parallel deliver 900 W to IP switches and routers**

Upgrade and modernization

In the case of modernization, most data centers are reaching limits in terms of increased power per board. Chassis and cabling is already installed, power supplies and power backup are often limited to what is already in place, and refrigeration systems have limited room for extra cooling. This means that data center operators will have to select a critical part of the system for upgrade, especially when new cards might require 1.2 to 1.5 kW each, which is almost twice that of today's average power per board.

Most present boards use two 450 W quarter-brick modules in parallel (Figure 1), delivering an average power of 900 W (actually, about 800W when thermal derating is taken into account). For next-generation boards, designers usually opt for higher-power quarter-bricks modules because they must use to existing footprints. These offer up to 860 W in their datacom version (45 V to 60 V input), Figure 2, or as stacked modules at the same power, but for wider input range of telecom standards (36V to 60V/75V), Figure 3. These upgrades do not offer the most efficient architecture when for demands per board in the 3-to-5 kW range.

The problem goes beyond supply

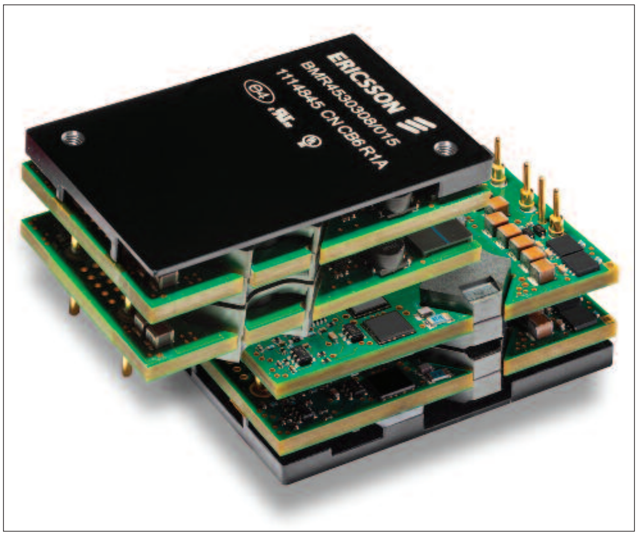

LEFT Figure 3: **Wide-input-range telecom stacked power modules**

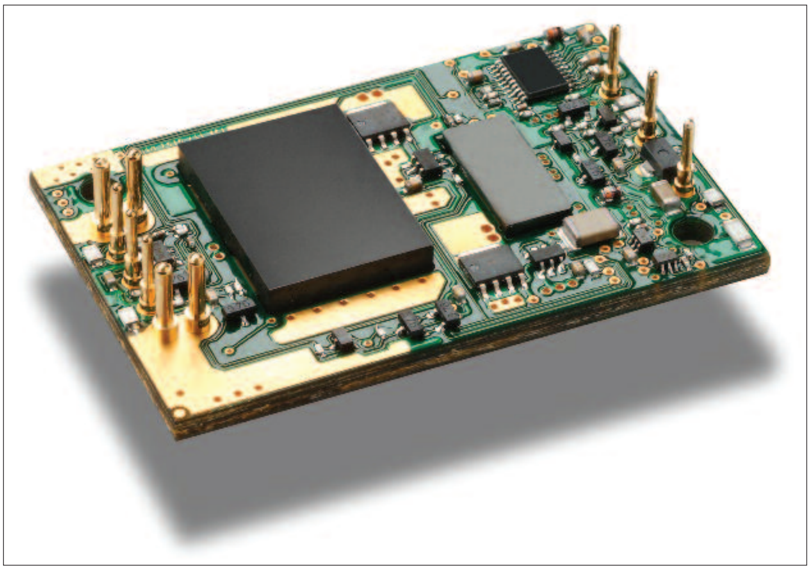

BELOW Figure 4: **The Double-P high-power-density module, launched in 2003**

circuitry. PC boards have become thicker, with heavy copper layers to carry the 100 A of board-mounted power sources (1200 W/12 V) to the loads; at the same time, the board may have up to 30 layers.

Therefore, soldering of power modules to the boards – and rework in the case of defects – has become a critical factor. Board power architects have also started to question use of just a single pin per voltage rail, to deliver the power to the intermediate bus.

These issues led to a new power-brick concept in 2003, with the 'The Double-P' (Thermally Enhanced Double Power Pin) initiative which featured two power pins and a layout to optimize power dissipation through pins and ferrites (Figure 4).

Along with delivering power comes the cooling problem and conventional air circulation may no longer be adequate. Solutions include use of heat-pipes to move heat from processors and other dissipating components to cold plates and heatsinks. In these systems, the board-power module baseplates are also connected to master cooling elements, reducing thermal stresses and improving reliability. These techniques are acceptable when upgrading existing installations, but will likely not be suitable for the multi-kilowatt boards of the next generation.

New developments and data centers

As demands increase, system architects will have to reconsider overall routing, shortening connections and increasing the number of I/O points per board in order to limit interconnection losses which introduce performance-killing latencies. This will require bigger boards plus cooling via refrigerated cold-walls for greater efficiency than conventional ventilation.

Architects are also re-thinking the paralleling the quarter-brick unit, the packaging of which has not evolved significantly since high-density versions were introduced in the late 1990s. In addition to the difficulty of soldering bricks to boards with up 45 layers, the outputs of these power bricks must deliver current to the multiple rails to distribute the 250 A (and later, 415 A). As a result, architects may demand products optimized for cold-wall conditions, resulting in a new type of board-power source which actually takes up more space, but has improved thermal exchange, interconnections enhanced for the higher currents and new assembling techniques.

The final product could be similar to the 1.5 kW module (Figure 5), with modules including larger transformers with interconnection to high-power, press-fit sockets pre-installed during board preparation.

Conclusion

When the 'The Double-P' initiative started, 800 W boards were challenging enough yet 10 years later, boards are at 1.2 kW and quarter-bricks are nearly reaching a kilowatt. Achieving 3 kW/board soon and 5 kW/board by 2020 will almost certainly require new technologies which combine unprecedented levels of innovation across multiple dimensions.

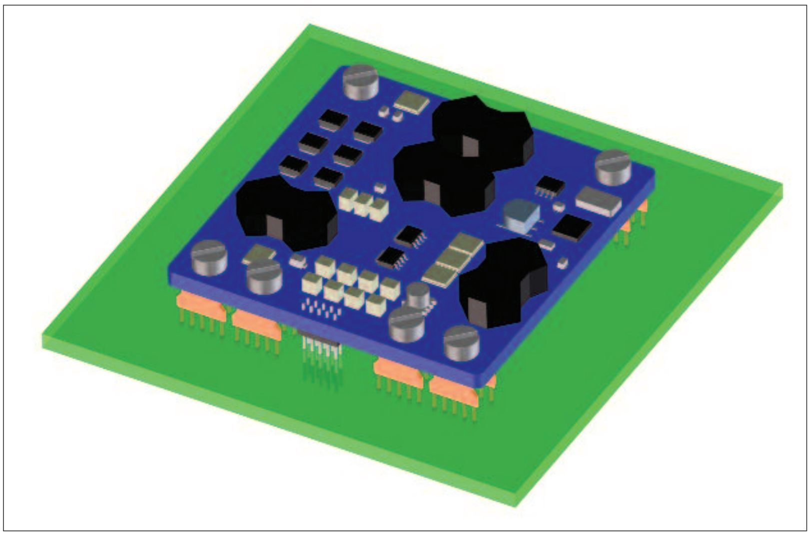

Figure 5: **Simulation of 1.5 kW isolated power modules designed for efficient cooling and easy assembling required to power multi-kilowatt boards**