IBM's Next Generation: Servers Look to Bus Converters

by Dr. Ray Ridley

There is no doubt that there is an urgent push to commoditize power supplies. And in some high-volume areas, commodity prices are being achieved. Many millions are made for a given specification, but often at the expense of performance. This is not true advancement toward commodity products. It is simply aggressive cost reduction in light of large production quantities.

What the industry has been searching for is a solution that can be applied to a wide array of power levels and applications. For this, the solution of outsourcing and a blind eye to compromised performance will never be acceptable. A true revolution is needed in the way power systems are structured to change today's design practices. Only if we arrive at standardized building blocks, whether they are circuit elements, or actual products, can we ever hope to achieve fast and reliable design. In this futuristic scenario, applying a succession of appropriate building blocks to a board would equal a complete power supply.

One attempt to achieve this is the bus converter technology - power building blocks that are standardized for mass production, but which perform just a reduced subset of the total system power requirement. In last year's issues of SPM, we showcased products from two companies-Vicor and IR-that achieved astonishing power density and efficiency by reducing the tasks required of a power block.

If we are ever to reach the point in this industry where we can engineer good power supplies from building blocks, each of which provide part of the function of the total power system, we need dramatic improvements in performance.   What is needed in a building block to make this work is:

- High Density
- High Efficiency
- High Reliability
- Low Cost

Many research groups have tried in the past to define power supply building blocks, but always failed in one or more of these categories. Several issues ago, we talked about the new Vicor building blocks based on data sheet information. We also showed how IR could achieve high densities with discrete components and conventional SMT packaging on FR-4 boards. These are, if anything, sophisticated building blocks for the experienced designer.

Using a set of customized building blocks, IBM is incorporating them in a most surprising place. According to Kevin Covi, a Senior Technical Staff Member at IBM Corporation in Poughkeepsie, New York, the building block approach will be used as part of their 20 - 30 kW power system for high-powered servers. This caught us by surprise since the available building blocks are only about 200 W each. We expected to hear of more companies working at lower power levels. However, in thinking further about it, this perhaps makes sense. At the 200 W power level and below, there are multiple options for low-cost power. Also, the competition is fierce with the brick converter market, and switching frequencies using almost conventional PWM topologies are already quite high.

At the multi-kW power levels, packaging and transformer technologies are large and expensive for high current outputs. Switching frequencies are limited by circuit parasitics to just a few hundred kHz. Replacing these large power supplies with multiple smaller supplies allows the switching frequency to be raised into the MHz frequencies. This can lead to dramatically smaller assemblies.

Another surprise in IBM's application is that they are presently processing power a total of three times between input power and output load. Once, for power factor correction; once for isolation; and once for load regulation. They have demonstrated that multiple power processing is actually more efficient that trying to do every function in one single stage. The new system will add a *fourth* stage of processing - with concurrent improvements in performance.

IBM's Three-Phase Power System

The very first issue of Switching Power, in July 2000, presented details of IBM's three-phase power system. (You can find details of this at www.switchingpowermagazine.com) This power system is probably the best thought-through, best-implemented, and reliable, high-power computer system in the world today.

Since the server power consumption is so high - in the order of 20 - 30 kW per frame, three-phase power feed is essential. Power factor correction, and low line distortion is also mandated. This is achieved with the novel approach of using three single-phase power factor correction circuits across each line-to-line voltage. To handle the worldwide range of voltages from 208 - 480, a buck + boost topology is used, the details of which can be found in our first magazine issue.

The outputs of each of the single-phase PFC units cannot be directly tied together. Fig 1 shows how DC-DC isolation stages allow a regulated 350 VDC from each phase to be connected in parallel to provide a high voltage bus capable of delivering 20 kW. IBM is unique in choosing to use this isolated high-voltage bus for power distribution around the computer cabinet. Most companies use 48 V, but this necessitates the use of heavy copper wire or bus bars to distribute the necessary power. With 350 VDC, small-gauge wires can be used to efficiently distribute the power.

When this system was first designed several years ago, processor logic voltages dominated at 1.6 to 3.3 V. And 200 - 400 A load converters were effective at delivering power to the loads. Now the processor voltages are moving down towards 1 V, total system currents are higher, and the architecture of Figure 1 is pushed to the limits to efficiently deliver power. Like all other designers at these voltages, IBM replaced output rectifier diodes with synchronous rectifiers to increase efficiency, but the output modules are bulky and difficult to place in close enough proximity to the load.

Furthermore, processor chips often require customized voltages to optimize speed and efficiency. (That is the reason for the programmable outputs of VRMs.) The high-current bulk power approach of Figure 1 will not be able to provide the flexibility and packaging density for the next generation servers.

The front end of the IBM power system would be extraordinarily hard to improve upon. It is reliable, rugged, and covers exceptionally wide range input with no changes to the circuit-all while maintaining efficiency in excess of 90%. The challenge for the next generation IBM system, therefore, lies in the 350 VDC to load conversion.

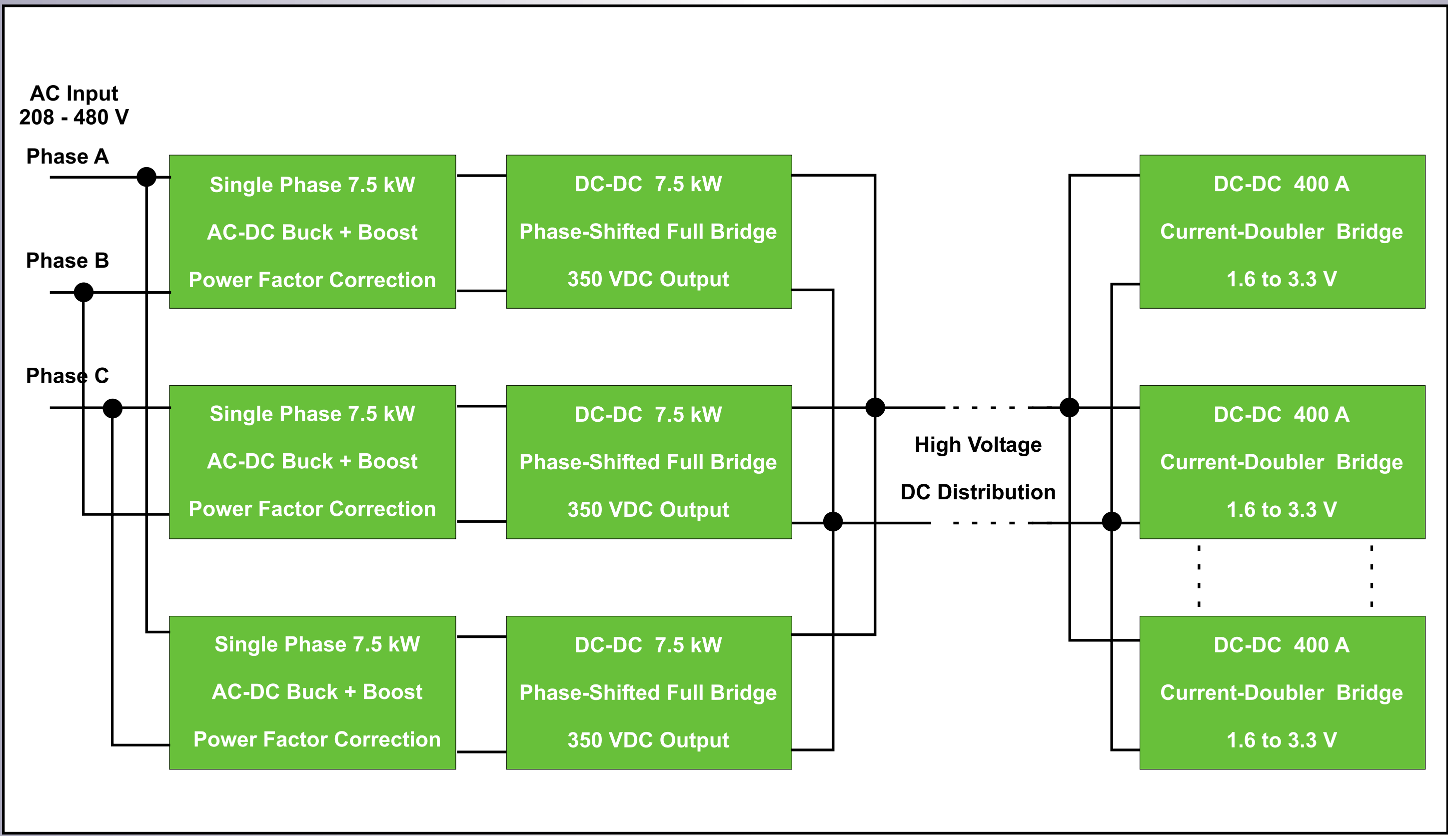

Fig. 1: Present Architecture of IBM Server Power System

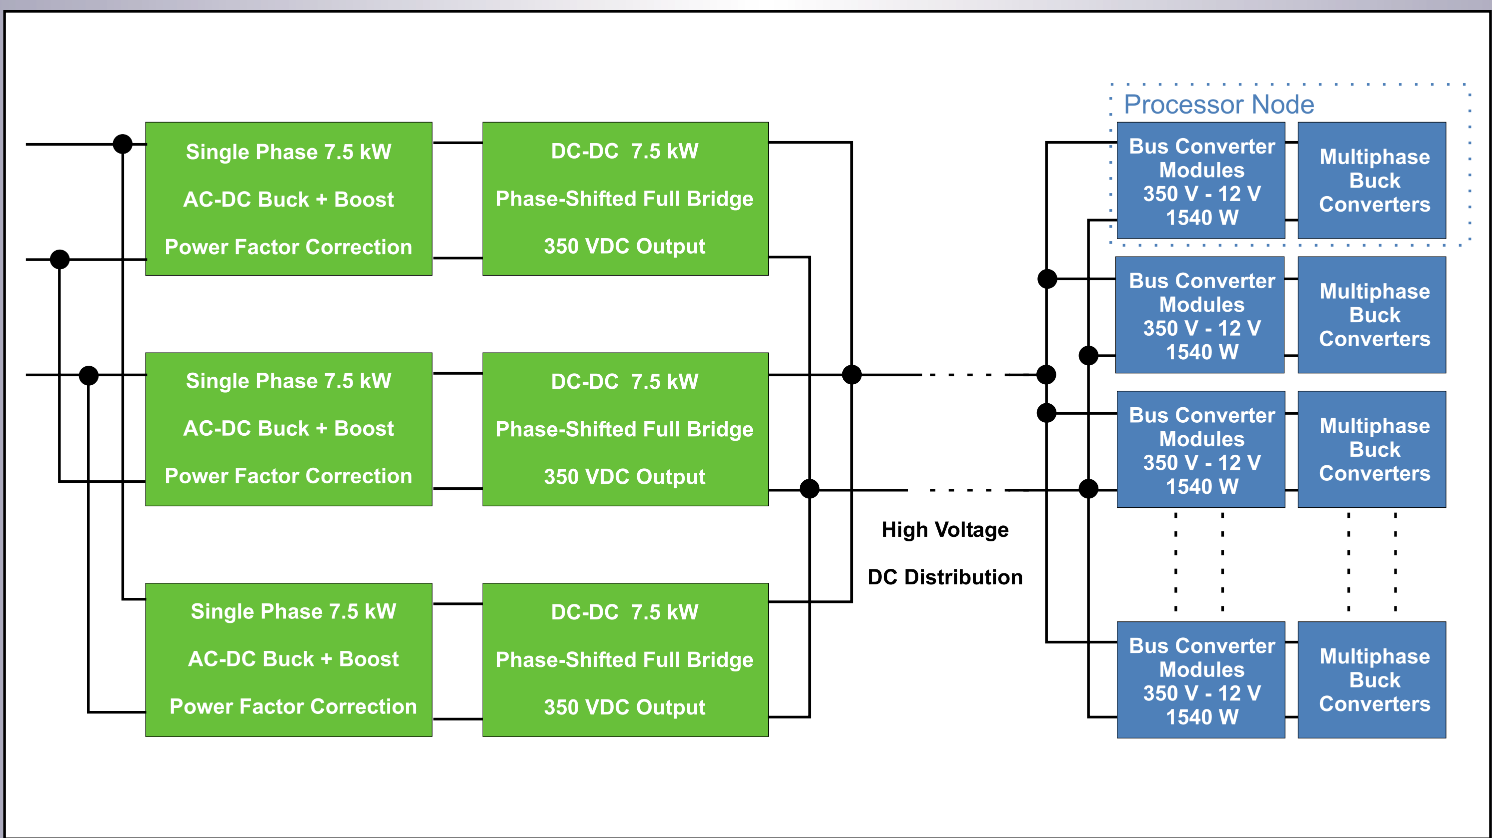

Fig. 2: Proposed Next-Generation Architecture of IBM Server Power System

The requirements for the next generation final power conversion from 350 V to processor voltages are as follows:

- Increased power density
- Increased granularity of output voltages for individual programming
- Increased reliability of overall power system
- Fast design turnaround.

IBM's new system will actually add a fourth power conversion stage into the mix. The final stage of power conversion will use VRM technology with multiphase buck converters. VRM technology has made amazing advances in the last 5 years, and multiphase buck converters can provide high current outputs from 12 V inputs with outstanding performance.

However, IBM's power bus is at 350 V, not 12 V. Power distribution of 20 kW would be completely unfeasible with 12 V distribution. There are three possible solutions-designing a bulk 350 to 12 V converter, designing new generation 350 to 1 V converters, or purchasing a pre-designed module for 350 to 12 V conversion.

The proposed final approach is shown in Figure 2. Fixed conversion-ratio modules, with a custom step-down ratio of 350 V down to 12 V will be used. These modules are currently under development at Vicor, and IBM expects to have samples within the next month. This increased input voltage will actually raise the power of the module by an additional 20% to 240 W over the current 48 V version of the modules. The modules inherently current share when hooked together, and are easy to synchronize, making them ideal for high power levels.

As we will see in the next section, the present 48-V units from Vicor are very close to achieving all the goals needed for the system.

Building Block Noise Issues

As we have mentioned in SPM previously, the specifications from Vicor for power density and efficiency are far beyond the industry norm, and we have confirmed their performance in our own labs. Of utmost concern when purchasing pre-packaged modules is noise performance, and reliability. Far too often, we have been called in to

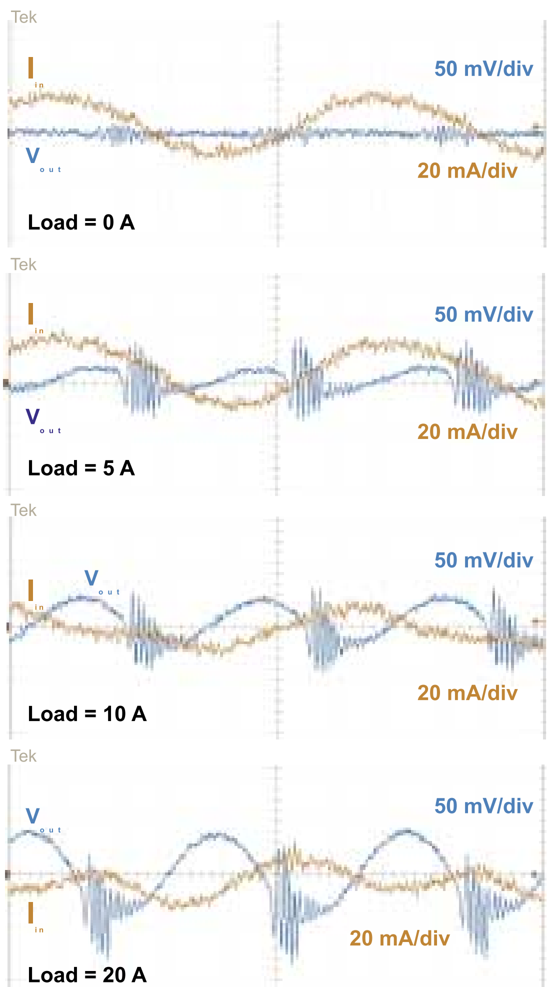

Fig. 3: Noise Levels of the Vicor 48:12 Bus Converter Module at Different Loads

trouble-shoot power systems composed of purchased modules that simply don't work well in the system due to the amount of noise they produce.
In one extreme case, noise was appearing on the back-plane of a telecomm system, and was traced back to the case of a converter that exhibited 100 V peak-to-peak noise at 1 MHz! Needless to say, this radiated very effectively throughout the equipment, causing many problems with data transfer errors.

In other situations, we have found common mode noise to be excessive, and have been unable to suppress it due to a lack of accessibility to the noise-generating nodes of the internal circuitry. And we have observed that as densities rise, in many cases the common mode noise increases. Planar transformers increase primary to secondary capacitance, and harmonic-rich squarewave switching drives large currents into the capacitance, resulting in high common-mode noise.
In contrast, the Vicor modules are one technology that promise and deliver outstanding noise performance. Zero-current, and zero-voltage switching remove fast edges that cause problems.

Figure 3 shows the output voltage noise and input current noise, measured with the 48:12 Vicor BCM. Note: unlike many other promised technologies, these noise levels are with NO additional output filtering, and just a small capacitor on the input to filter the input current. The noise at full power is just 120 mV peak-to-peak, as shown in the lower traces of Figure 3. There are two dominant components, one at 3.5 MHz, the output ripple due to the internal ripple frequency, and 200 MHz.

Figure 4 compares the Vicor noise levels to that of other technologies. In Figure 4a, you can see the classic noise spectrum of a 100 kHz PWM converter. A triangle wave is generated at the switching frequency, but the noise is dominated by spikes (caused by the square-wave switching). It is very hard to effectively filter out these noise spikes due to the rapid slew rates in the converter.
The noise of the Vicor converter is relatively narrow-band, and is easy to filter so that it is not a problem in your power system. We can't emphasize enough how important this is - so often, a converter is packaged very densely, but the required filter components greatly reduce the final power density, and can cause tremendous engineering problems late in the power system design cycle.

Figure 4c shows the output voltage noise of a discrete power supply offering similar power density to the Vicor converter. This is a demonstration converter from International Rectifier. While the IR converter matches efficiency, and offers flexibility of design, it does not provide the required noise performance for most systems.

Figure 5 shows an all important measurement indicating common-mode noise performance. The scope waveforms are made with the ground and the tip of the oscilloscope probe both connected to the return of the power supply output. This is not, of course, a calibrated setup, but the amount of noise measured with this configuration gives a good indication of how much trouble you will have with a converter when incorporating it into your power system.

Differential-Mode Noise levels of the Vicor 48:12 Bus Converter Module Compared to other

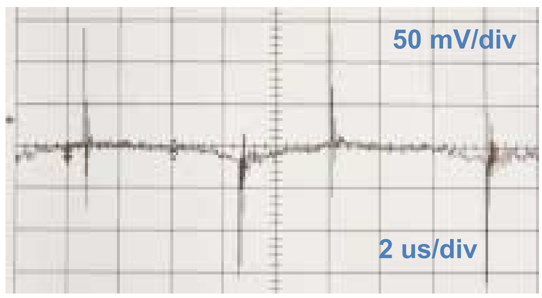

4(a) Conventional 100 kHz PWM Forward

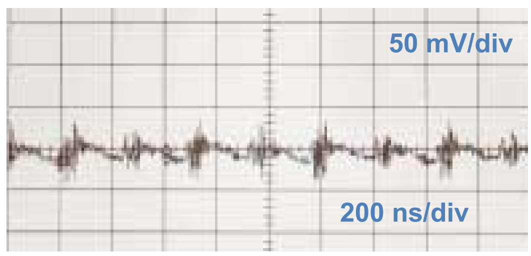

4(b) Vicor 3.5 MHz Bus Converter

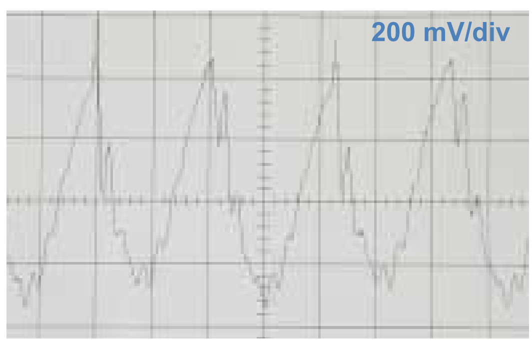

4(c) IR's Half Bridge Bus Converter
(Note different scale)

A conventional PWM converter again shows the classic noise waveform - switching frequency triangular ripple, with high-frequency switching spikes dominating the waveform. The peak-to-peak amplitude in this measurement is 130 mV.

In contrast, the Vicor Bus Converter generates only about 30 mV peak-to-peak restricted bandwidth noise. This will be easy to reduce with additional filtering.

The final open question is-is it practical? When it comes to cost and reliability, the greatly-reduced parts count of the Vicor BCM, and published reliability figures suggest that this will not be a driving factor. And the cost will no doubt be greatly driven by how widely this kind of approach is accepted by our industry.

Further System Testing

We are anxiously awaiting further test results from IBM for their power system. If the 350 V input modules meet the same specifications of power density and efficiency that have been demonstrated at 48 V, we will have a powerful new tool in our design arsenal for future power

Common-Mode Noise Vicor 48:12 Bus Converter Module Compared to PWM Supply

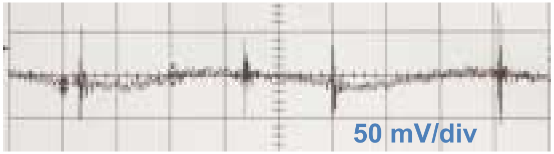

5(a) 100 kHz PWM Forward

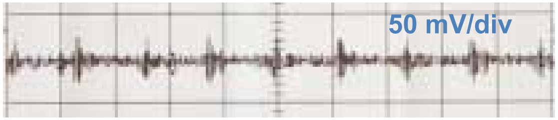

5(b) 3.5 MHz Vicor BCM

supplies. It will then be possible to do three-stage power conversion (PFC, Bus Conversion, Point-of-Load Conversion) that will exceed the performance of conventional power supply design.

While the 48 V converters are fun to contemplate, they really offer no useful solution for offline power. IBM's push for higher-voltage bus converter modules will hopefully change things for the better.