The quandaries of memory subsystem design

Designing a memory subsystem always seems to come with trade-offs – depending on how you use your system, you usually have to give up one benefit to get a more critical one. For instance, graphics systems typically place the emphasis on speed. For a desktop application, you need to ensure that parts are inexpensive and readily available. For a server, one of the main objectives is to improve performance by maximizing the amount of memory.

Servers pose a particular dilemma in that you can't simply add large amounts of memory to get better performance – the maximum amount of memory has traditionally been constrained by the existing system architecture. But what if you could have it all in a server – a memory solution that gave you increased performance AND higher density, yet is still cost-effective. Would you be interested?

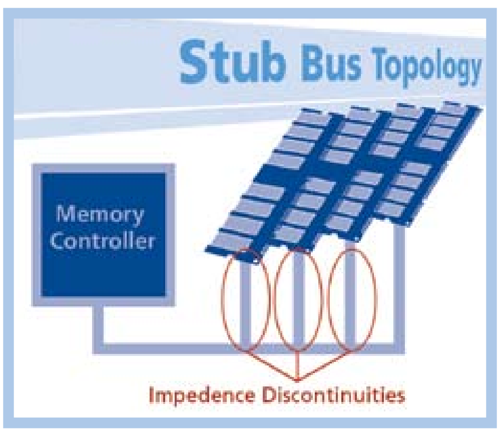

Fig. 1*: Existing stub bus topology -- modules are connected to memory controller in parallel.*

An overview of current server memory technology

For some time, the memory solution of choice for servers has been registered DIMMs (RDIMMs).

RDIMMs allow servers to utilize higher-density modules than desktops. These modules often have 18 or 36 DRAM chips on them (compared to 8 or 16 on a typical desktop solution). In order to handle the increased loading on the chipset, register(s) on the module introduces a delay that effectively adds an additional clock cycle to the CAS latency (CL). This tends to limit performance, especially when data is read from non-consecutive address locations.

Additionally, RDIMMs connect to the memory controller in parallel; that means each module must be directly connected to the controller. In order to accomplish this, server designers connect the memory to a "stub bus," a single electrical connection with branches that connect to each memory module. With as many as 72 or more connections in today's systems, the signal quality may degrade where the bus and DRAM devices meet, potentially causing errors, especially as speeds increase.

Because of these factors, systems using RDIMMs can utilize only a limited number of memory modules. As DRAM clock/data rates increase, the number of devices allowed on each memory channel decreases because of the impedance discontinuities associated with the stub bus topology.

Until now, designers had only two choices: limit the amount of memory so that fewer errors occur at higher speeds, or accept slower speed to gain the density required. But when capacity, reliability, and performance are critical, neither of these choices will do.

This limitation presents even more of a challenge to the industry when you consider that server performance is driving memory capacity to double approximately every two years.

A brighter future for server memory

So we need more memory, but can't just add more DIMMs. So how do we solve the problem?

It seems logical that the first option would be to increase the density of the DRAM components more quickly. If only it were that easy! As densities increase, die sizes increase, as well. Additionally, the highest-density DRAM usually has a price premium compared to mainstream densities – this is why today you usually see a 2GB registered DIMM made with 512MB DRAM instead of 1GB or 2GB DRAM.

Another option would be to add additional memory channels. A downside to this is that higher-pin-count controllers aren't necessarily feasible in many cases, and in the cases where it is possible, it adds complexity, takes up space, and adds cost to the system. Again, this potentially becomes cost-prohibitive in many systems.

Enter the fully buffered DIMM (FBDIMM). FBDIMMs feature an advanced memory buffer (AMB) that buffers the DRAM data pins from the channel and uses high-speed point-to-point lines (serial signaling similar to PCI Express) to eliminate the stub bus. This bus architecture allows modules to be connected in series.

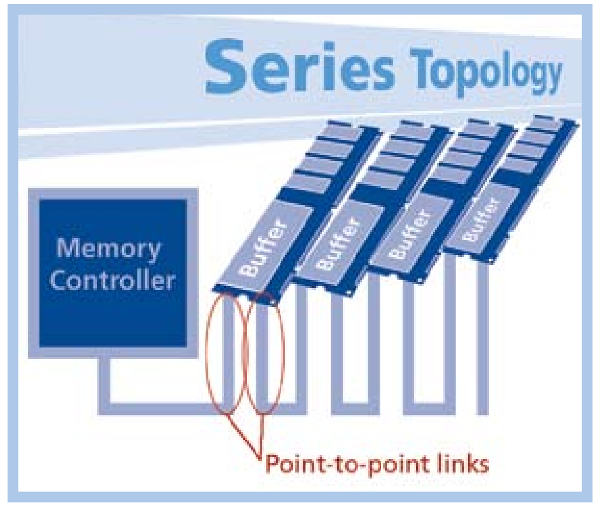

Fig. 2*: FBDIMM series topology -- modules are connected to memory controller in a series.*

As a result, server designers can dramatically increase the number of modules in a server while also increasing speed and reliability.

FBDIMMs offer a relative ease of design compared to a wide parallel bus and also offer improved signal integrity. In addition to allowing for simpler motherboard routing, FBDIMMs help to eliminate the need for stacked modules. Stacked modules traditionally have been used to maximize the amount of memory in a system. Since the memory controller can support more modules, it's no longer necessary to stack DRAM on modules.

An additional benefit of FBDIMMs is that they help simplify memory transitions (for example, DDR2 to DDR3) due to similar sockets and topologies. This structure offers an additional advantage (cost and time savings) as systems migrate to new memory technologies, making FBDIMMs attractive in telecommunications and networking applications.

The serial data path of FBDIMMs also allows FBDIMMs to achieve faster performance than RDIMMs. Their simplified connection structure means FBDIMMs reduce the cost and complexity of system design while offering faster transmission rates than conventional architecture. Plus, they can perform reads and writes simultaneously, while eliminating the read-to-read delay between data transfers. The peak throughput is about 1.5 times that of a single 72-bit DRAM channel.

Dramatically improved error detection

FBDIMMs enable unprecedented reliability, serviceability, and accessibility (RAS) features at the module level. First, FBDIMMs incorporate a cyclic redundancy check (CRC) that provides greater data and address/command protection than RDIMMs. For even greater defense against errors, FBDIMMs also utilize a bit-lane fail-over correction feature that identifies bad data paths and removes them from the operation. They also take advantage of new error-detection methods to reduce address/command soft errors and are hot-swappable. Combined, these error-detection methods dramatically reduce errors in server memory.

Mechanically, a FBDIMM is similar to the RDIMMs we're used to seeing. However, there are a couple of differences. The first thing you'll notice is that FBDIMMs have a heat spreader on them. If you remove the heat spreader, you'll notice that the registers and PLLs on the module have been replaced by the AMB. They also are keyed differently than DDR2 RDIMMs (need a different socket) to prevent a FBDIMM being placed in a RDIMM slot and vice-versa. FBDIMMs are roughly 30mm tall, the same height as DDR2 RDIMMs. A 38mm version is also in development to allow for 8GB and 16GB modules, without stacking.

Fig. 3*: FBDIMM with and without heat spreader.*

Proof that the FBDIMM solution works

So does the improved architecture of FBDIMMs really improve server performance? Consider the following comparisons of RDIMM-based systems vs. FBDIMM-based systems configured at three levels.

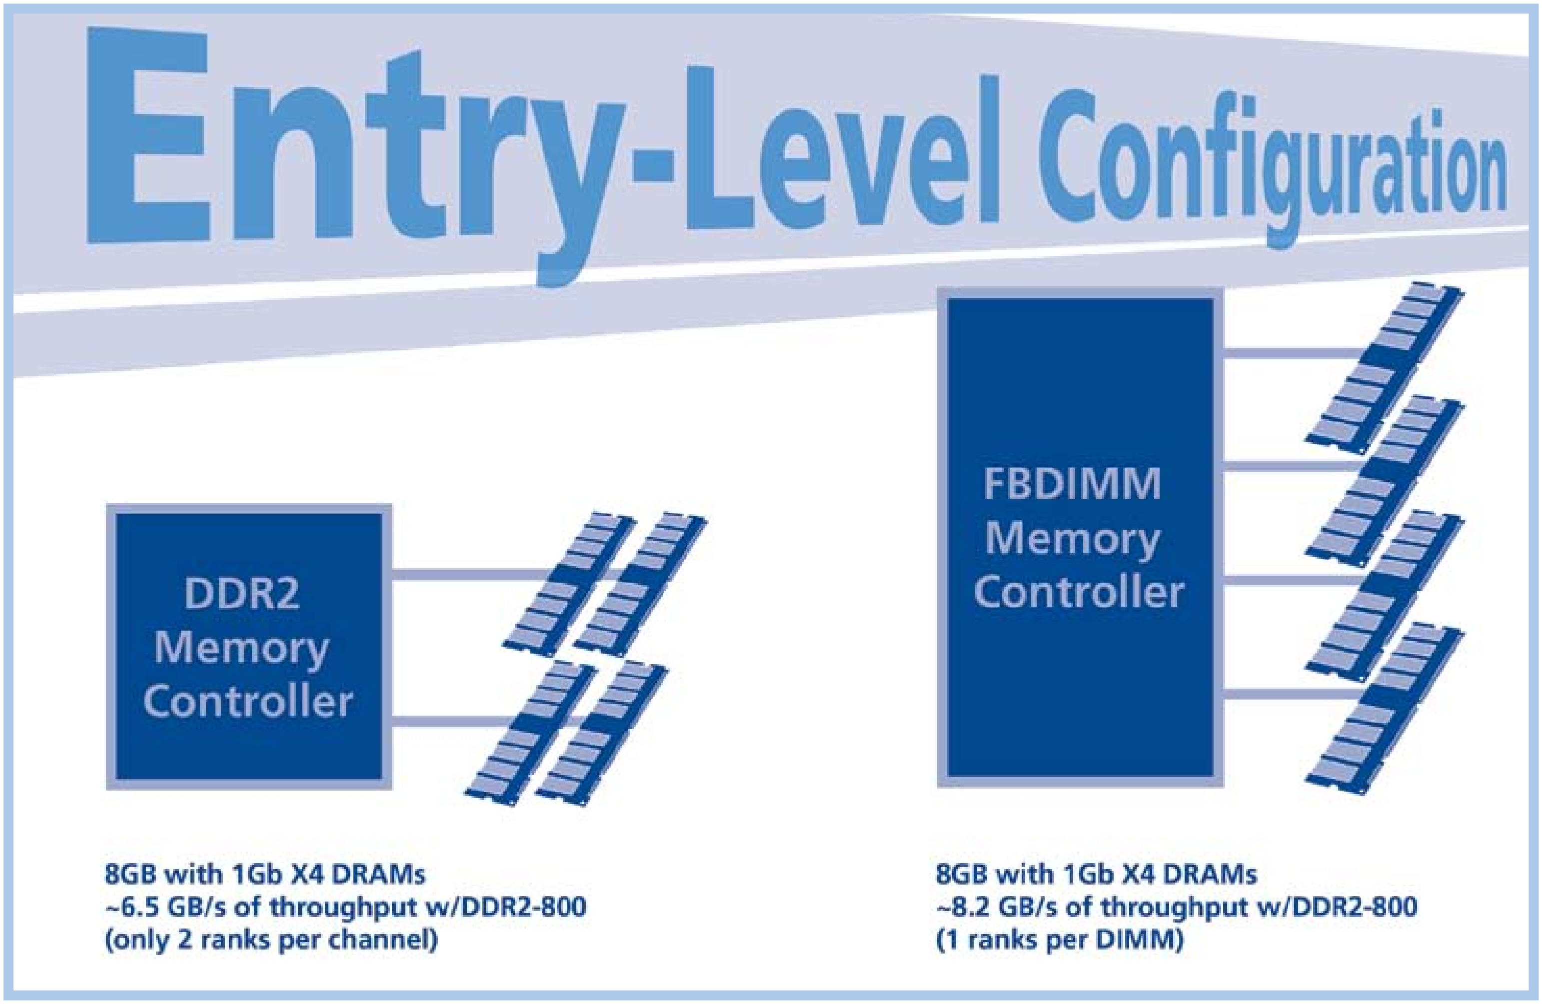

- **Equal capacity 8GB vs. 8GB**
- **Better throughput ~6.5GB/s vs. ~8.2GB/s**
- **Lower pin count ~480 vs. ~420**

In an entry-level system, you have the same amount of memory as a DDR2 RDIMM-based one. However, there are two significant performance improvements. First, the bandwidth is higher. Second, the number of pins required by the memory controller/motherboard is significantly reduced, which is a great benefit to system designers and users. Having fewer pins simplifies the system design and helps to reduce system cost.

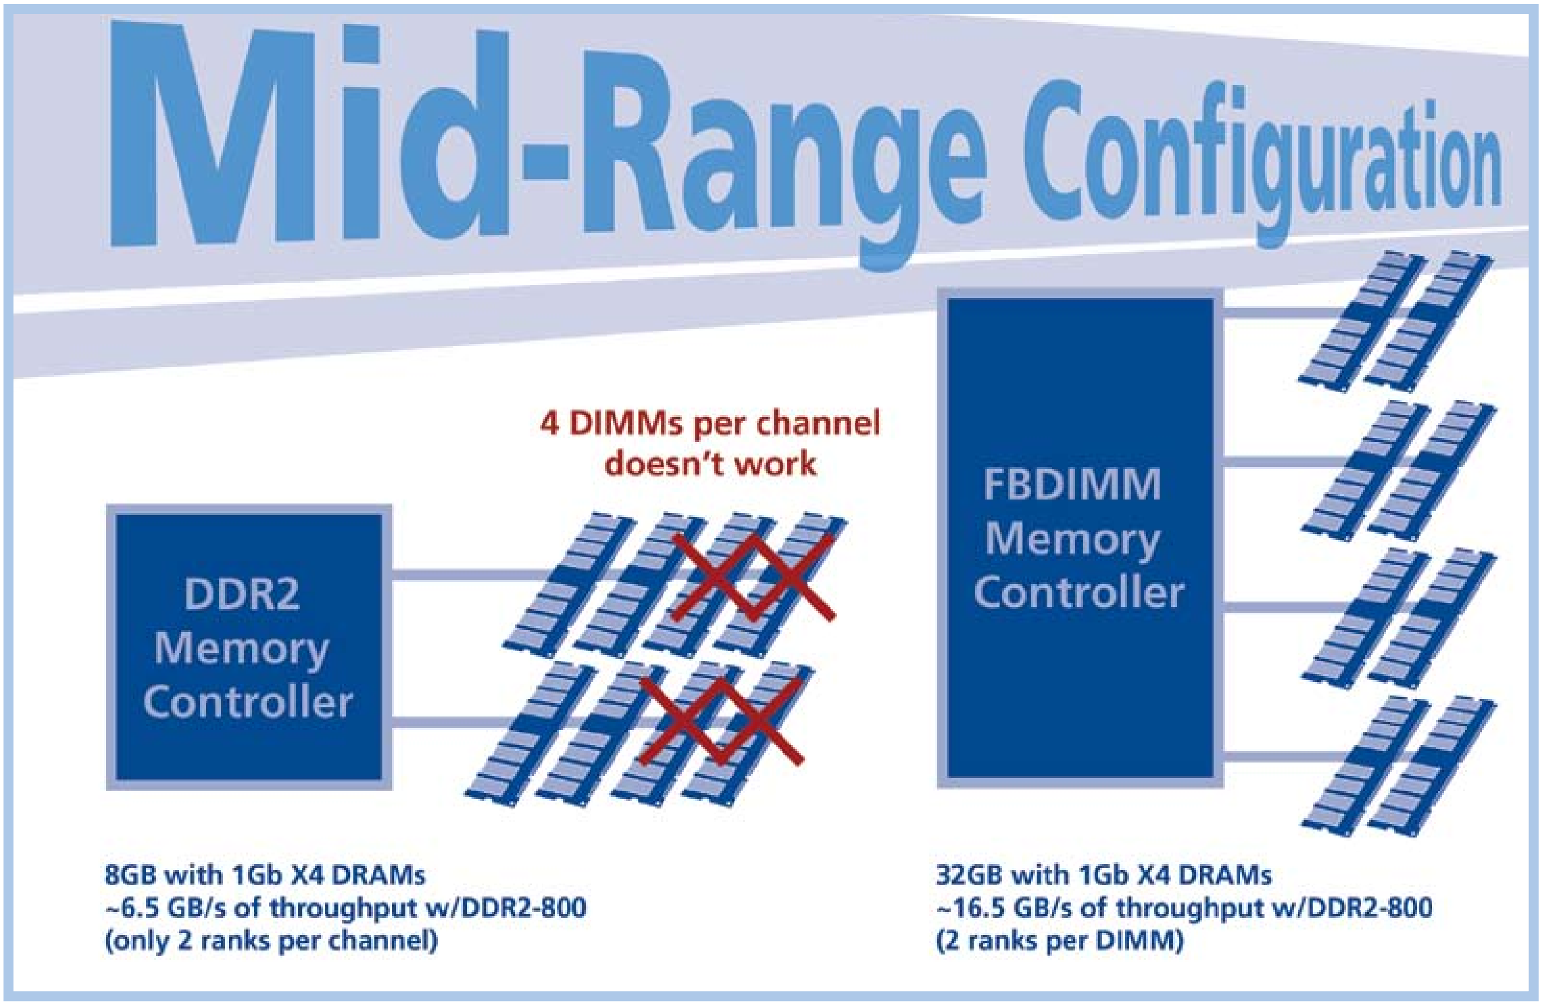

- **4X capacity 8GB vs. 32GB**
- **2.5X throughput ~6.5GB/s vs. ~16.5GB/s**
- **Lower pin count ~480 vs. ~280**

In a mid-range RDIMM system, you're hard pressed to increase the amount of memory. At 533MHz and beyond, the stub bus architecture makes it difficult to maximize the amount of memory, and you really can't support additional DIMMs per channel. You're already seeing systems where, in order to maximize performance (in other words, maximize memory), you have to utilize single-ranked 2GB DIMMs, which have a cost premium compared to dual-rank 2GB DIMMs.

A mid-range FBDIMM based system eliminates this issue and offers significant additional capacity and bandwidth.

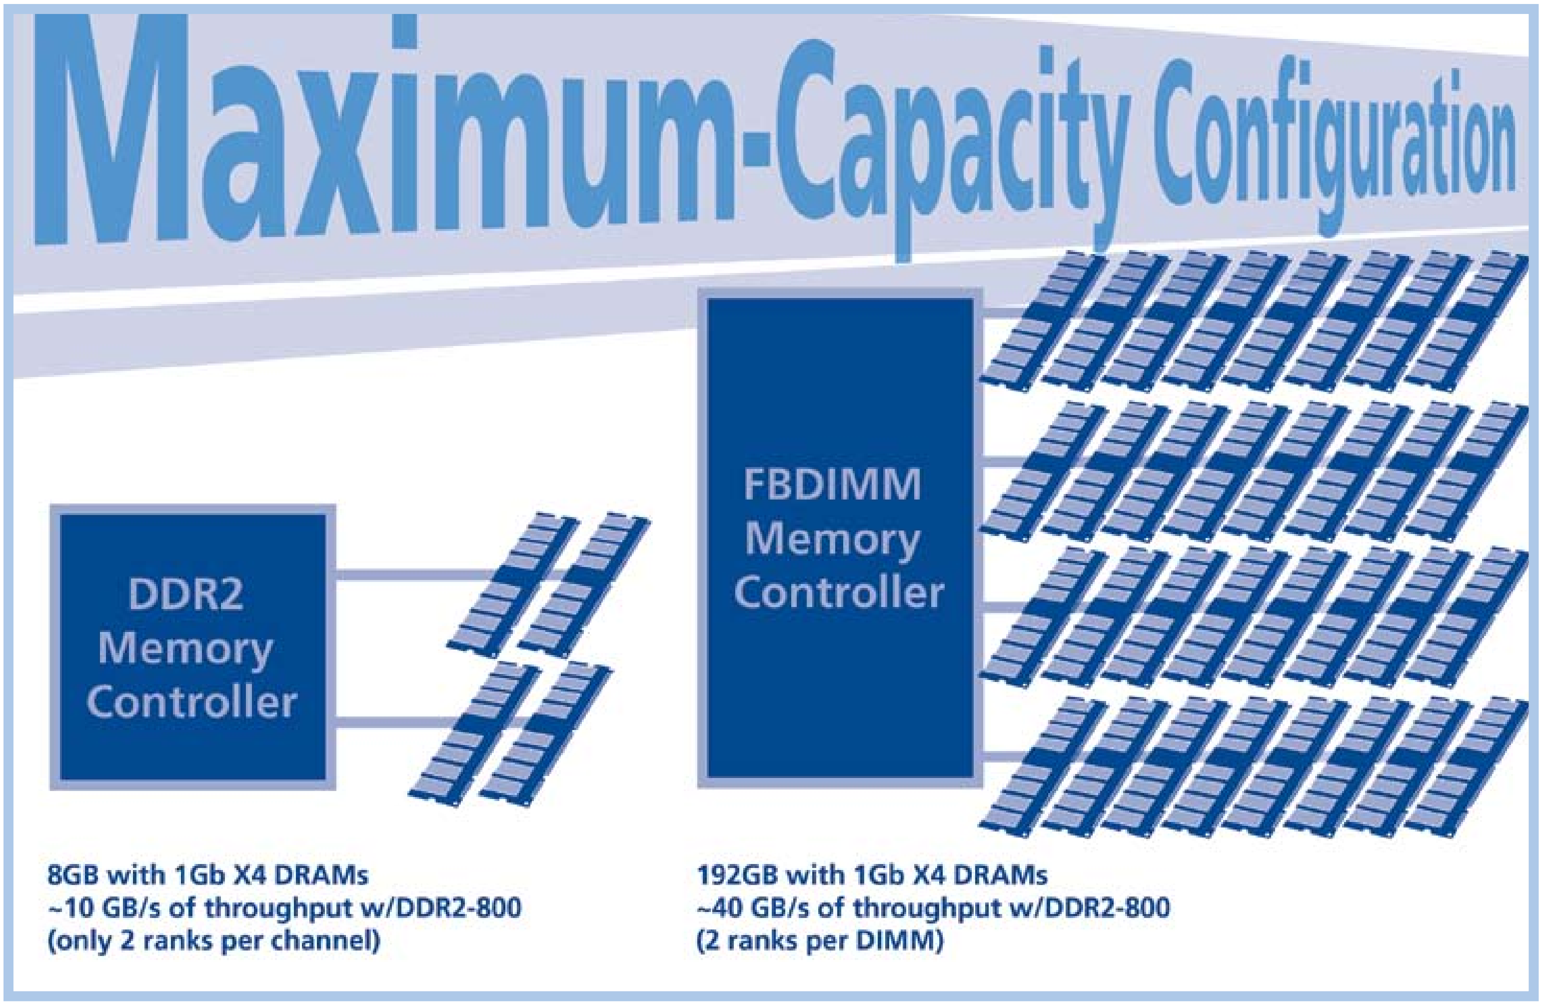

- **24X capacity 8GB vs. 192GB**
- **4X bandwidth ~10GB/s vs. ~40GB/s**
- **Lower pin count ~480 vs. ~420**

In a high-end system, many of the system limits that currently exist with registered DIMMs are removed. Basically, you're limited only by the number of DIMM slots you can populate. In this configuration, you can have 24X the memory with roughly 4 times as much bandwidth!

Where to use FBDIMMs

FBDIMMs are ideal for mission-critical servers that require maximum performance and a low error rate.
Eventually you'll see them used in workstations and high-end desktops as well, as you can have more than one module per channel as data rates increase. Future applications also include networking and telecommunications equipment.

Just remember that FBDIMMs and RDIMMs are not compatible, so you can't replace current server modules with FBDIMMs.

FBDIMMs solve the server and workstation capacity problem. They offer density and performance improvements and unprecedented reliability, and the ease of design enables additional applications. You'll see early system introductions in mid-2006 for performance servers, based on the Intel® Bensley platform. Larger systems and low-end servers are expected to migrate to FBDIMM technology as higher data rates are required and/or as FBDIMMs become mainstream in 2007.

About the Author

Andy Heidelberg has roughly 20 years of experience in the electronics industry. He began his Micron career in 1995, building memory modules. Since 1999, he has been part of the Engineering team at Crucial Technology, a division of Micron.

The Crucial advantage:

- Manufacturer-direct availability
- Limited lifetime warranty
- Free, industry-leading technical support
- Competitive pricing

Crucial has FBDIMM parts...

Call a sales representative or place your order online!

In t he U.S. and Canada
Toll-free: 888-363-2566
Phone: 208-363-5790 (Outside the U.S. and Canada)
E-mail: cruccust@micron.com
www.crucial.com/fbdimm

In the UK and EU
Freephone: 0800 013 7394 (UK)
Phone: +44 (0)1355 586016 (International)
Fax: +44 (0)1355 586289
www.crucial.com/uk/fbdimm (UK)
www.crucial.com/eu/fbdimm (EU)

crucial™ TECHNOLOGY
A Division of Micron