David Watts
Robert Moon

System x3755 Technical Introduction

Introduction

Delivering an industry-leading, 64-bit framework for high performance computing and business computing, the IBM® System x3755 is built for speed. The x3755 eliminates system architecture bottlenecks through the use of separate, high-speed links between the processors, main memory, and I/O. The AMD Direct Connect Architecture and HyperTransport technology combined with IBM Xcelerated Memory Technology™ drive the x3755 to deliver the performance, availability, and manageability needed for the next generation of high performance computing servers that require low latency and high speed access to memory.

This paper describes the technical aspects of the server. The topics in the paper are:

- Overview of the x3755
- Current models
- The x3755 design
- Processors
- System memory
- Serial Attached SCSI
- Graphics subsystem
- PCI subsystem
- Redundant Cooling Kit
- Redundancy
- Light path diagnostics
- Baseboard Management service processor
- Remote Supervisor Adapter II SlimLine
- PowerExecutive
- Supported operating systems

Overview of the x3755

Figure 1 The IBM System x3755

The following are the key features of the x3755:

- Four-way capable server in a rack-dense 4U form factor.
- One or two standard AMD Opteron 8000 Series Dual Core Processor, upgradable to three-way or four-way. Processors support 64-bit addressing with the AMD64 architecture.
- Support for AMD Virtualization (AMD-V).
- 1 GB or 2 GB memory standard expandable to 128 GB (using 4 GB DIMMs), using high performance PC2-5300 CL5 ECC DDR2 RDIMMs.
- Xcelerated Memory Technology to reduce latency for large memory configuration and enable each of the possible 32 memory slots, to run at the higher clock speed of 667 MHz.
- Chipkill™ on 1, 2, and 4 GB DIMMs.
- Online spare memory.
- One full-length PCI Express x16 slot, two full-length PCI Express x8 slots; one full-length PCI Express x4 slot; two full-length 64-bit, PCI-X 1.0, 100/133 MHz slots; and one half-length HTX slot.
- Integrated Adaptec AIC-9580W serial-attached SCSI (SAS) controller with support for internal RAID arrays using RAID-0, RAID-1 and RAID-10. Additional support for RAID-1E, RAID-5 and RAID-6 using an optional ServeRAID™ 8k adapter. MegaRAID 8480 adapter is also supported for external SAS storage with the EXP3000 enclosure.
- Four internal hot-swap SAS drive bays.
- Integrated Dual-port Broadcom 5708C Gigabit Ethernet PCI Express with TCP/IP Offload Engine (TOE) and failover capabilities.
- Baseboard Management Controller standard with optional Remote Supervisor Adapter II SlimLine adapter.
- Three-year warranty, on-site, 9 hours per day, 5 days per week, with a next business day response.
- The x3755 is targeted at businesses requiring a rack-optimized server that is optimized for high performance computing due to the server's low latency and high speed access to memory. Applications include weather simulation, CAE, crash analysis, financial analysis, fluid dynamics, protein folding, and the like.

Figure 2 shows the front panel of the x3755 showing the four hot-swap drive bays and the DVD-ROM drive.

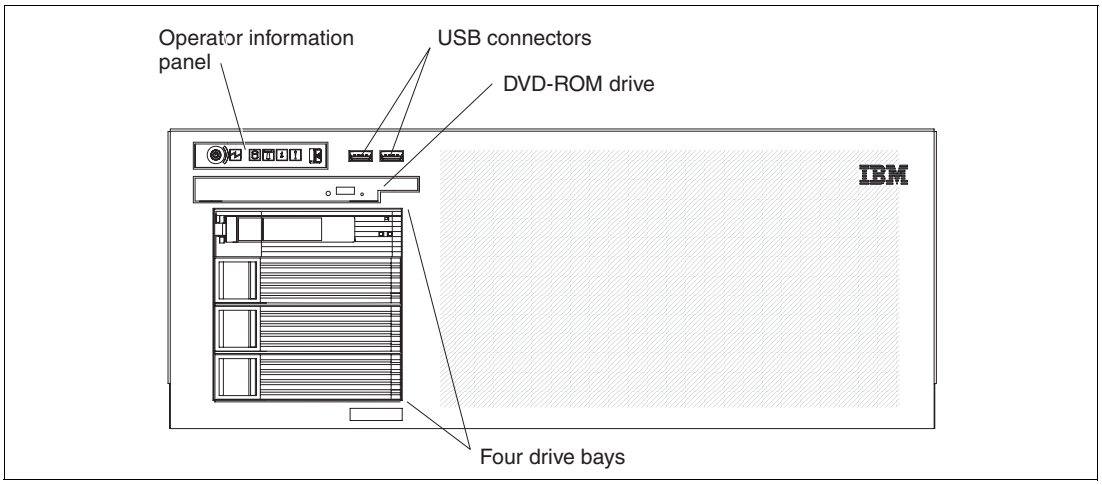

Figure 2 Front panel of the x3755

Figure 3 shows the rear connectors.

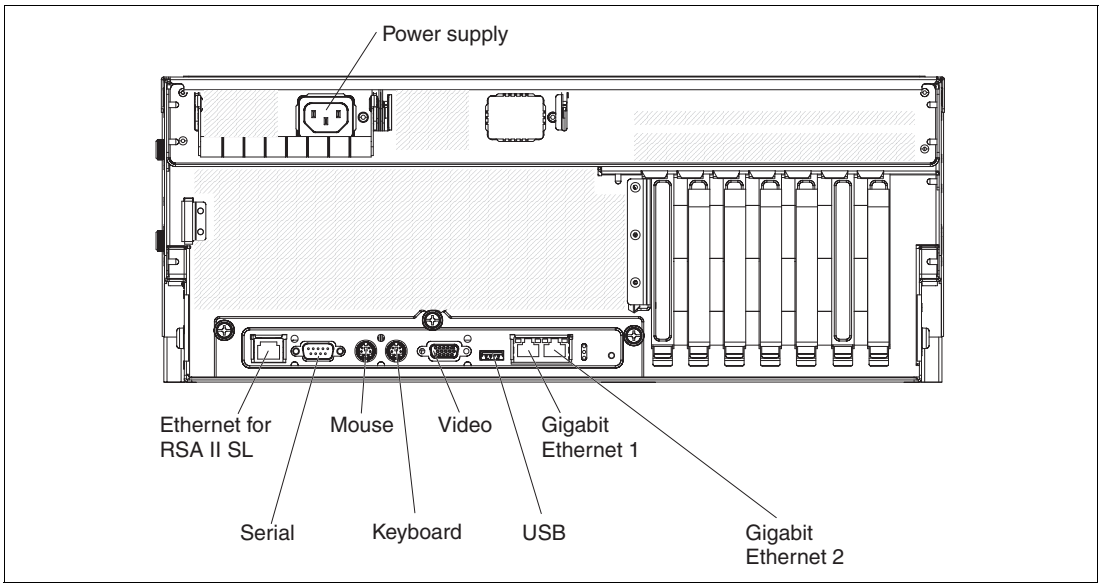

Figure 3 Rear view of the x3755

Current models

The following x3755 model was announced in August 2007.

Table 1 x3755 models announced in August 2007

Model	CPU standard	L2 cache	Memory std	CPU/mem cards	PassThru cards
8877-7RU	2x 3.2 GHz Opteron 8224 SE	1+1 MB	4 GB (4x 1 GB)	2 / 4	0

The following x3755 models were announced in April 2007.

Table 2 x3755 models announced in April 2007

Model	CPU standard	L2 cache	Memory std	CPU/mem cards	PassThru cards
8877-3RU	1x 2.4 GHz Opteron 8216	1+1 MB	1 GB (2x 512 MB)	1 / 4	1
8877-5SU	2x 2.8 GHz Opteron 8220	1+1 MB	4 GB (4x 1 GB)	2 / 4	0
8877-6RU	2x 3.0 GHz Opteron 8222 SE	1+1 MB	4 GB (4x 1 GB)	2 / 4	0

The following x3755 models were announced in August 2006.

Table 3 x3755 models announced in August 2006

Model	CPU standard	L2 cache	Memory std	CPU/mem cards	PassThru cards
8877-1RU	1x 2.0 GHz Opteron 8212	1+1 MB	1 GB (2x 512 MB)	1 / 4	1
8877-2RU	1x 2.2 GHz Opteron 8214	1+1 MB	1 GB (2x 512 MB)	1 / 4	1
8877-4RU	2x 2.6 GHz Opteron 8218	1+1 MB	4 GB (4x 1 GB)	2 / 4	0
8877-5RU	2x 2.8 GHz Opteron 8220 SE	1+1 MB	4 GB (4x 1 GB)	2 / 4	0

Note: The U in the model numbers is for countries in North and South America. For EMEA, substitute G (for example, 1RG). For Asia-Pacific countries, the letter varies from country to country. Consult the announcement letter or the *IBM System x Configuration and Option Guide* found at:

http://www.ibm.com/support/docview.wss?rs=1201&uid=psg1SCOD-3ZVQ5W

The x3755 supports one, two, three, or four processors. A PassThru Card is required to support a configuration of either 1 or 3 processors. See “Processors” on page 11 for more information.

All models support a maximum of 128 GB using 4 GB DIMMs in 32 sockets. To achieve the maximum, you must install three additional CPU/memory cards (one or two CPU cards are standard), remove the standard DIMMs, and insert 32 4 GB DIMMs.

The x3755 models listed in Table 3 have one PCI Express x16 slot, two PCI Express x8 slots, two 64-bit 100/133 MHz PCI-X 1.0 PCI slots, and one HTX slot.

Express models: Certain countries also may offer preconfigured systems called Express models. These are preconfigured with more processors, memory, disk, and a ServeRAID controller. The advantage of these models is fewer part numbers to order.

The x3755 design

The x3755 uses a combination of AMD architecture and a ServerWorks chipset coupled with an IBM design. The architecture consists of the following components:

- One to four AMD Opteron processors
- One to four CPU/memory cards
- One PassThru Card in a one-way or three-way CPU configuration
- One ServerWorks HT2100 A PCI Express Bridge, one HT2100 B PCI Express Bridge, and one ServerWorks HT1000 South Bridge

Figure 4 shows the block diagram of the x3755 machine type 8877.

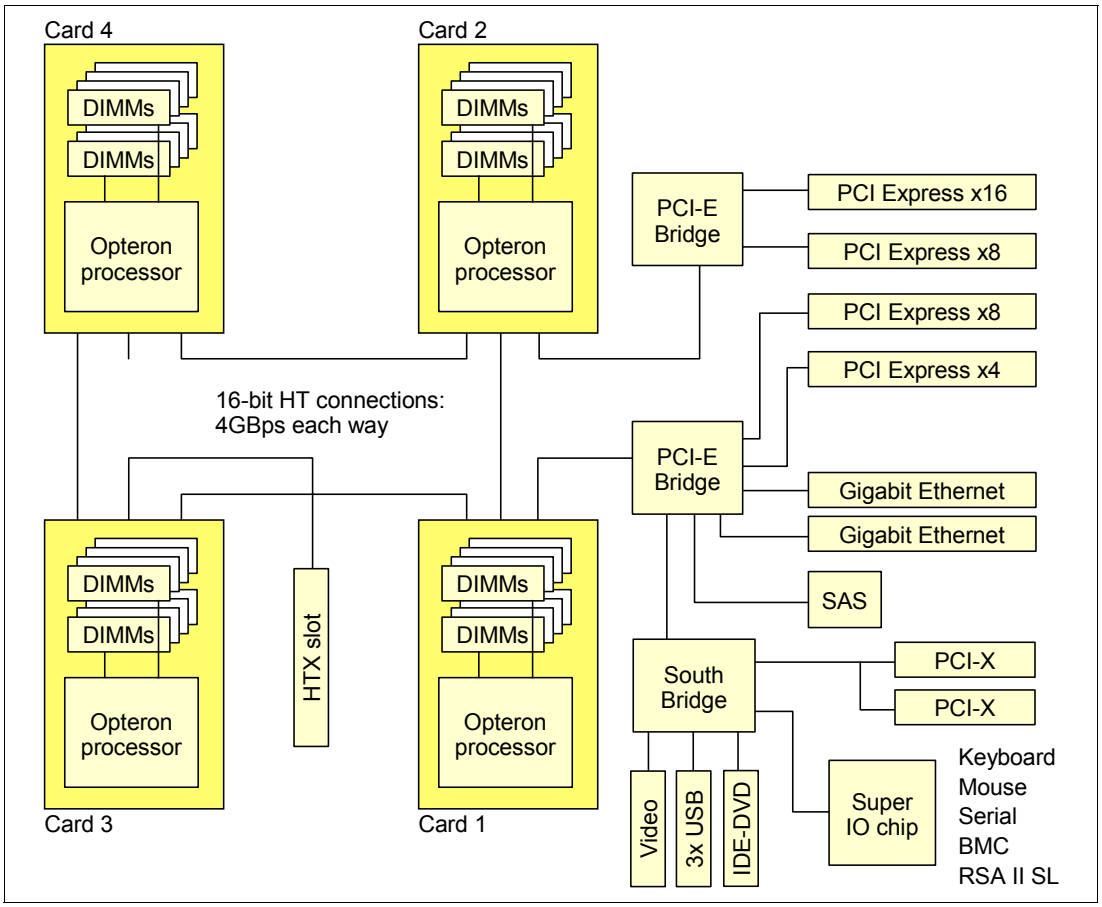

Figure 4 System x3755 8877 system block diagram with four CPU/memory cards installed

The x3755 has two PCI Express bridges. One provides one PCI Express x16 slot and one x8 slot. The second provides the other PCI Express x8 slot and one x4 slot, in addition to the onboard Ethernet controllers and the SAS onboard controller with optional ServeRAID 8k.

The south bridge provides the two 64-bit 100/133 MHz PCI-X slots, in addition to all of the onboard PCI and (super) I/O devices, including the Remote Supervisor Adapter II Slimline daughter card, if installed.

AMD architecture

The AMD Opteron processors do not use the typical shared front-side bus that is connected to a memory controller used in Intel®-based servers. Each Opteron processor has its own integrated memory controller and pins on the processor chip to directly connect to a memory bus. So, in Opteron, processor and memory controller logic are integrated into the same piece of silicon, eliminating the need for a separate memory controller part. Hardware vendors simply add a memory bus and memory DIMMs, and they have the core CPU and memory interface.

To keep data coherent between multiple Opteron processors, AMD introduced a new system bus architecture called *HyperTransport*. Three HyperTransport links are available on each Opteron processor, two used for CPU-CPU connectivity and one used for I/O. The two HyperTransport links used for CPUs enable the direct connection of two processors and the indirect connection of four or more processors.

With a four-processor configuration, the processors are placed at the corners of a square, with each line that makes up the square representing a HyperTransport connection between the processors. See Figure 5.

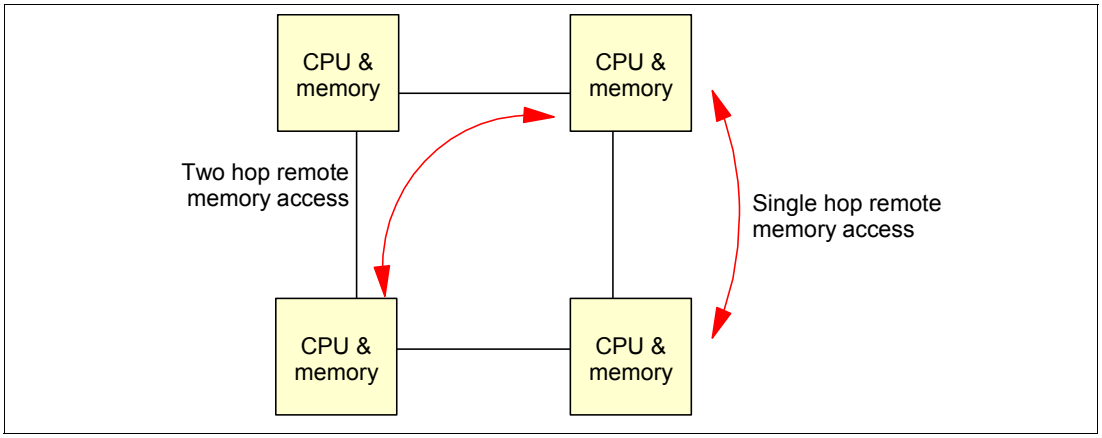

Figure 5 Remote memory access

With this design, whenever two processors on the same side of the square share data, the information passes directly over the HyperTransport interconnect to the other processor.

When this remote access occurs, it is called a *single hop remote memory access* and is slower than a local memory access.

However, when two processors on diagonal corners of the square share data or instructions, the information must travel through an additional processor connection before arriving at the diagonal processor. This extra hop adds overhead and is referred to as a *two hop remote access*.

The third port of the HyperTransport link is not used to interconnect the diagonal processors because it must be used for connection to a PCI I/O bridge, which connects such devices as PCI slots, network, disk storage, mouse, keyboard, video, and so forth.

The remote memory access latency of a processor accessing another processor's memory space makes the Opteron configuration a NUMA design. NUMA means that every processor has memory that is *closer* and thus more rapidly accessible and memory that is *remote* and slower that must be accessed via another Opteron processor.

AMD refers to its Opteron architecture as *sufficiently uniform memory organization* (SUMO) rather than NUMA. From an architectural standpoint, it still is a NUMA architecture, but the HyperTransport link is fast enough to run software written for SMP systems without very significant performance penalties. Current operating systems such as the latest versions of Linux® and Windows® 2003 SP1 support NUMA and make attempts to minimize remote memory transactions. But in practice, the percentage of remote memory accesses is largely determined by application behavior and by how data is manipulated by users of the application.

Figure 6 shows the Opteron architecture with the integrated memory controller and the HyperTransport connectors.

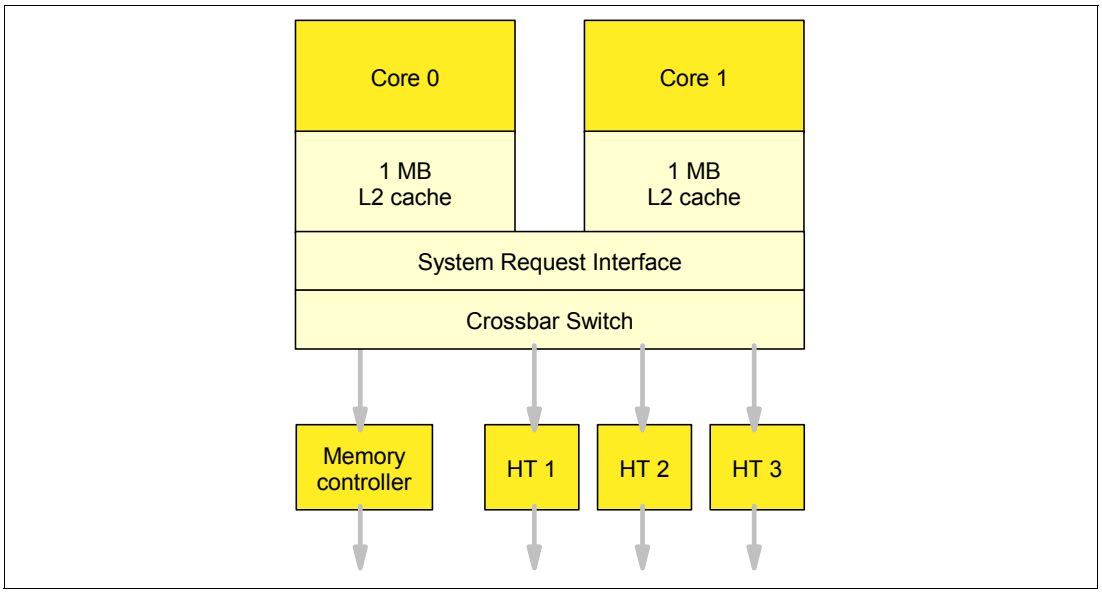

Figure 6 CPU architecture of the Opteron CPU with an integrated memory controller

HyperTransport

The HyperTransport architecture was initially developed by AMD but is now managed by an open consortium of several big IT companies such as AMD, Apple, Cisco, Broadcom, ATI, IBM, and many others. HyperTransport is an open standard for a high-speed, point-to-point link system that can be used for connecting a variety of chips.

The HyperTransport technology is used in devices such as network devices, graphics cards, or, as in the case of the AMD Opteron, as a high-speed interconnect for processors. The HyperTransport technology used for interconnecting Opteron processors is currently implemented at a speed of 1000 MHz with a bidirectional bandwidth of 4.0 GBps each way that leads to a peak full-duplex capacity of 8.0 GB per second per link. Current Opteron processors incorporate three HyperTransport links, which enables a peak bandwidth of 24 GBps per processor.

You can find more information about the HyperTransport and the HyperTransport Consortium at:

http://www.hypertransport.org/

IBM technology

With the x3755 IBM has introduced two new technologies:

- The Xcelerated Memory Technology
- The PassThru card

Xcelerated memory technology

The x3755 memory subsystem consists of two sets of four DIMMs in a daisy-chain interconnect (see Figure 7).

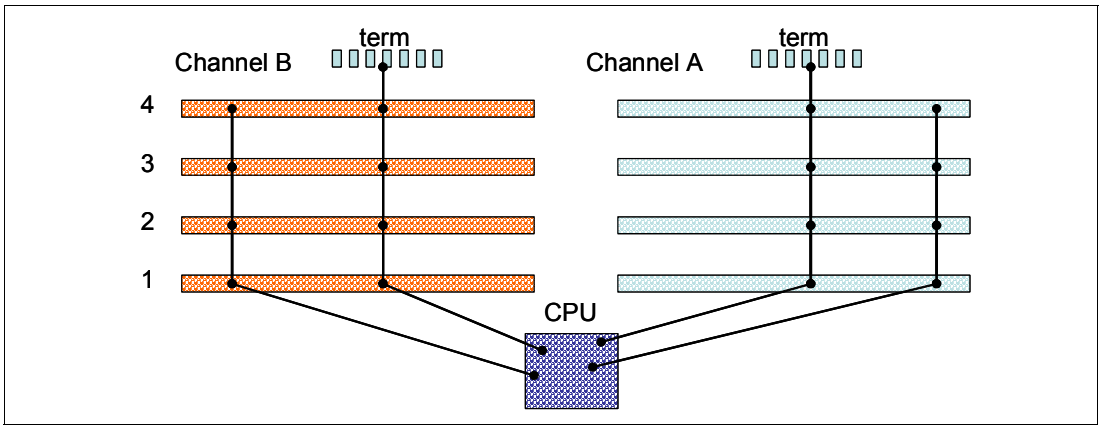

Figure 7 x3755 memory subsystem

Normally, if a read or write signal is sent down the memory bus to DIMM socket 4 (the DIMM furthest away from the CPU), then due to the design of the bus, this signal will be reflected by each DIMM along the bus. This creates additional signals that cause noise and can result in incorrect data reads and writes, which in turn could cause the system to hang.

The AMD design specifications say that if you add more than two DIMMs on a memory bus then you must lower the memory bus speed from 667 MHz to 533 MHz to minimize the effect of the noise. IBM developers, however, found that if you add a circuit to the bus that counteracts the noise, then you can maintain the timing and electrical integrity of the signal. This in turn means that you can keep the bus speed at the higher 667 MHz for all eight DIMMs on each CPU/memory card in the x3755.

This IBM unique design allows the x3755 to use the higher memory bus speed of 667 MHz even when using more than two memory DIMMs, thereby improving the memory performance of the x3755 versus the competition.

PassThru card

The PassThru card, part number 40K7547, is similar to the CPU/memory card except that it does not have any memory or CPU slots. The card was designed to address two issues with the AMD design in a one-way or three-way implementation.

- One-way configuration

  As shown in Figure 6 on page 7, each AMD CPU has three HyperTransport connections. One of these is used for I/O connectivity. Figure 4 on page 5 shows that the first bridge connects the master processor (card 1) to the majority of the I/O devices, including two of the PCI Express slots, the PCI-X slots, and storage. The second bridge connects a separate processor (card 2) to two additional PCI Express slots.

  To be able to have the second PCI bridge enabled, you are normally required to have a CPU/memory card installed. By introducing the PassThru card, however, IBM designers eliminated the need for and the associated cost of the second CPU/memory card. Figure 8 shows how the PassThru card is used in a 1-way configuration to enable the second PCI bridge.

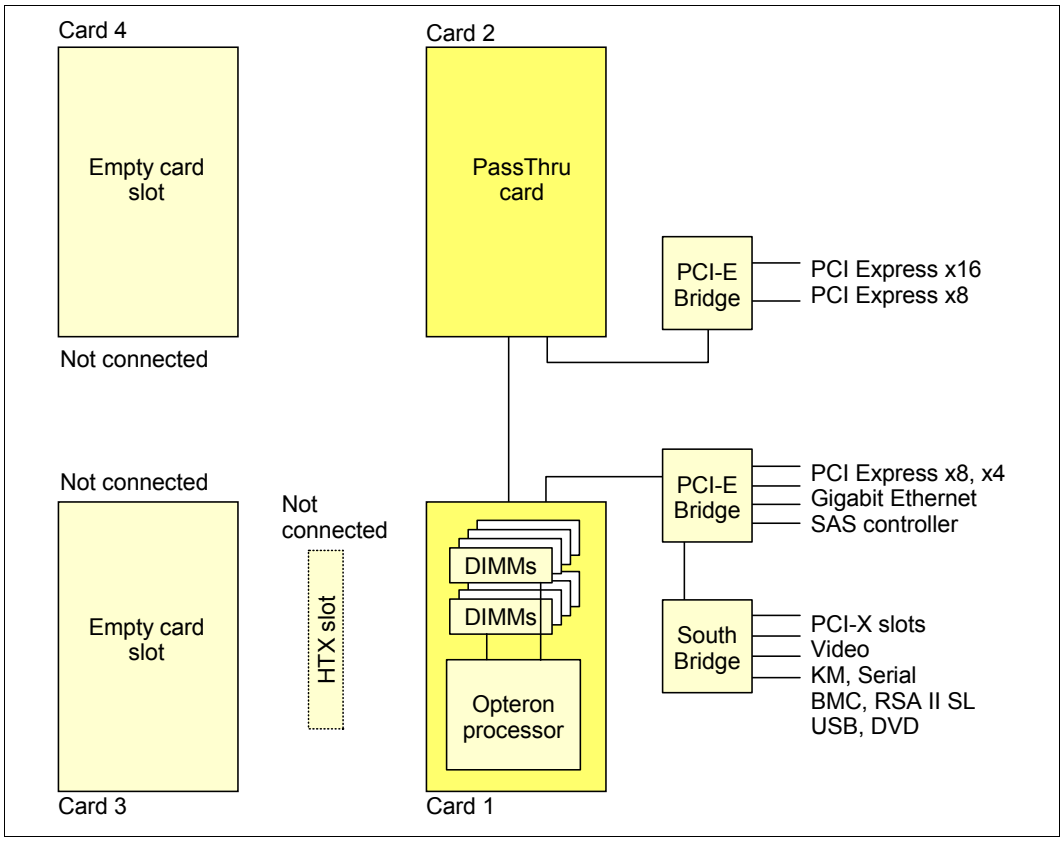

Figure 8 Block diagram of a one-way x3755 with PassThru card

- Three-way configuration

  As described in “AMD architecture” on page 5 and in Figure 5 on page 6, four processors are connected together by the HyperTransport links. These form a ring topology. This ring topology is broken when only three CPUs are used, as shown in the top half of Figure 9. The result is a bus rather than a ring. To prevent this from happening IBM utilizes the PassThru card, which can be inserted for the missing CPU, thereby re-enabling the ring topology, as shown in the bottom half of Figure 9.

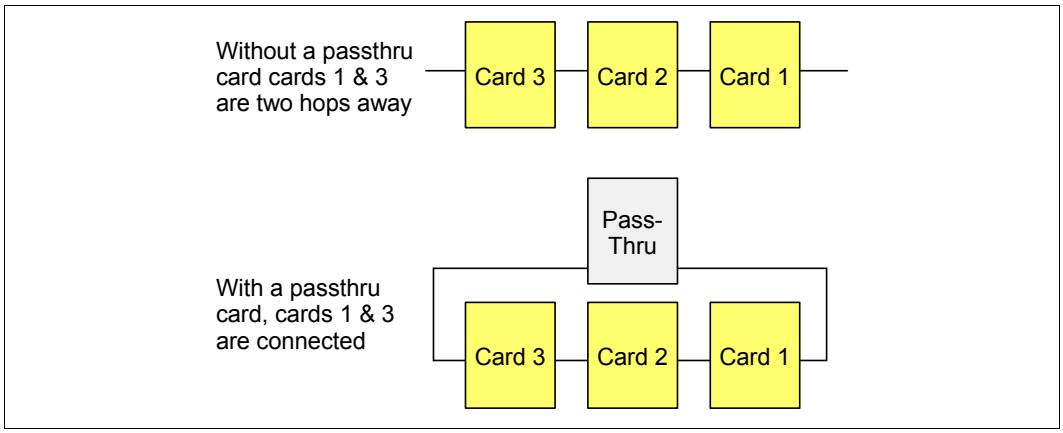

Figure 9 The benefit of the PassThru card for three-way configurations

Figure 10 shows a block diagram of the x3755 and how the PassThru card re-enables the ring topology.

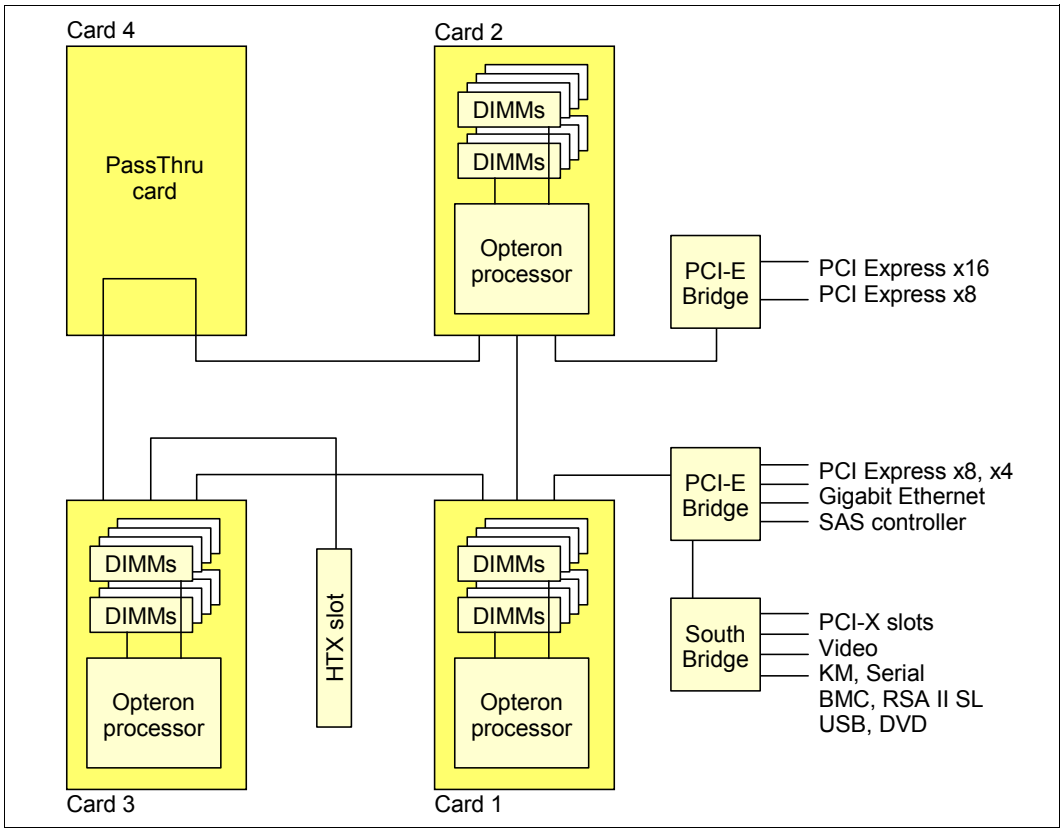

Figure 10 Block diagram of a three-way x3755 with a PassThru card installed

There are performance benefits to adding the PassThru card in a three-way configuration. Without the PassThru card, the configuration requires that snoop requests and responses originating from one of the two end processors and certain non-local references, to travel over two hops. With the PassThru card, this now becomes only one hop. The benefits in decreased latency and increased memory throughput are shown in Figure 11.

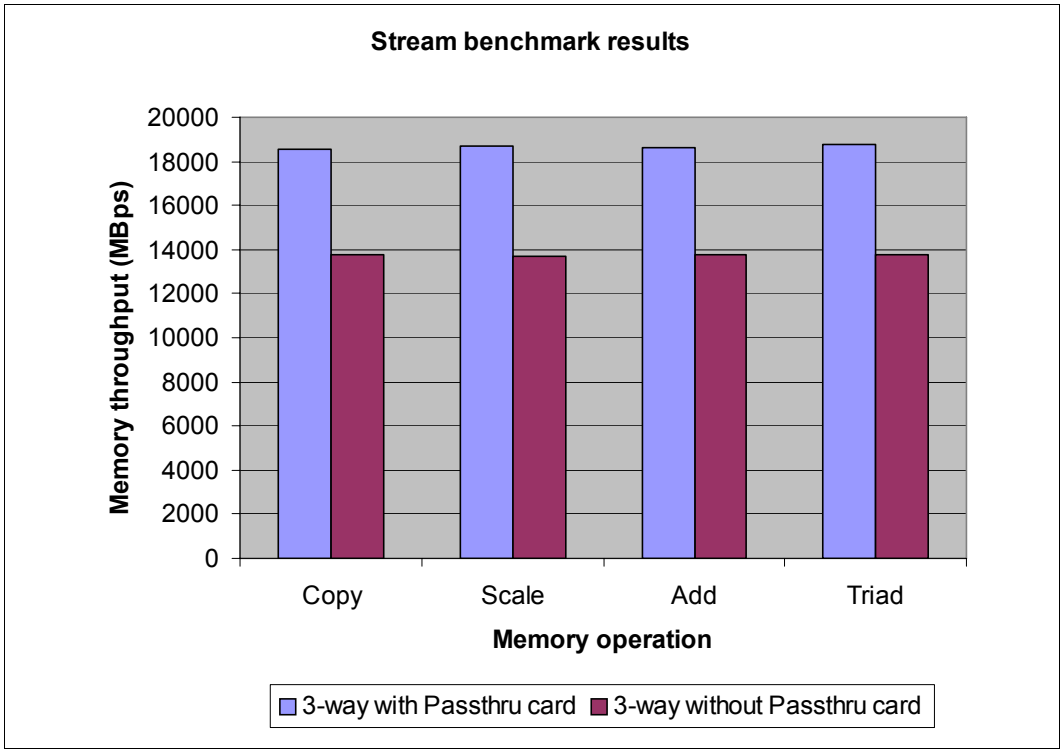

Figure 11 Memory throughput benefit of the PassThru card

For more information about this performance benefit see the white paper *Performance of the IBM System x 3755* by Douglas M Pase and Matthew A Eckl, available from:

http://www.ibm.com/servers/eserver/xseries/benchmarks/related.html

Processors

The x3755 models use the AMD Opteron revision F *Santa Rosa* processor with AMD64 extensions and AMD Virtualization (AMD-V).

AMD Revision F processors are the successor to the Rev E processors. They are not compatible with current Rev E processors but should be pin compatible with future quad core processors. The following features are incorporated into the Rev F processors:

- All rev F processors are multi-core. There will be no single core Opteron processors after the current generation.
- Support for DDR2 memory is added.
- The current 1 GHz HyperTransport technology is incorporated as the interface to initial Rev F processors.
- PCI Express support is added.
- AMD Pacifica virtualization technology and power management technologies are incorporated.

The major performance increase with Rev F AMD Opteron over the current Rev E AMD Opteron processors is related to the increase to DDR2 memory technology. DDR2 memory technology results in over a 25% increase in memory throughput for random memory accesses. For a performance analysis of DDR verses DDR2 see the IBM Redbooks® publication *Tuning IBM System x Servers for Performance*, SG24-5287.

Figure 12 shows the architecture of the Revision F processors.

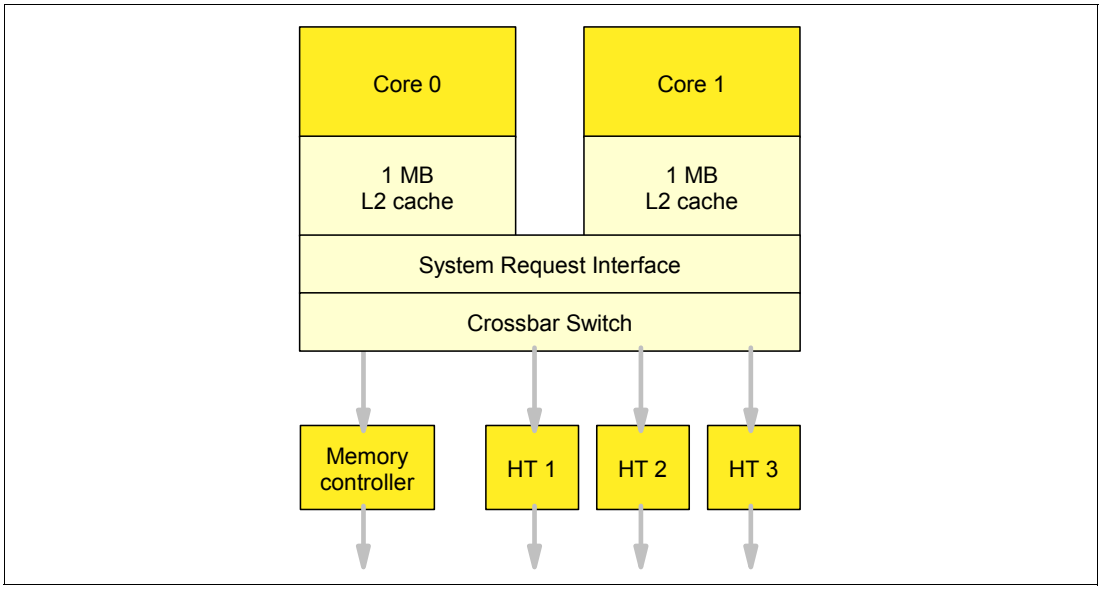

Figure 12 Architecture of dual core Opteron Rev F processors

Models of the x3755 come with one or two processors installed (see “Current models” on page 3). One, two, three, or four processors are supported.

Each processor is mounted on a CPU/memory card. When you order additional processors, the CPU/Memory card is included with the option, along with the CPU and the heatsink.

The CPU/memory card has a socket for the processor plus 8 DIMM slots as shown in Figure 13.

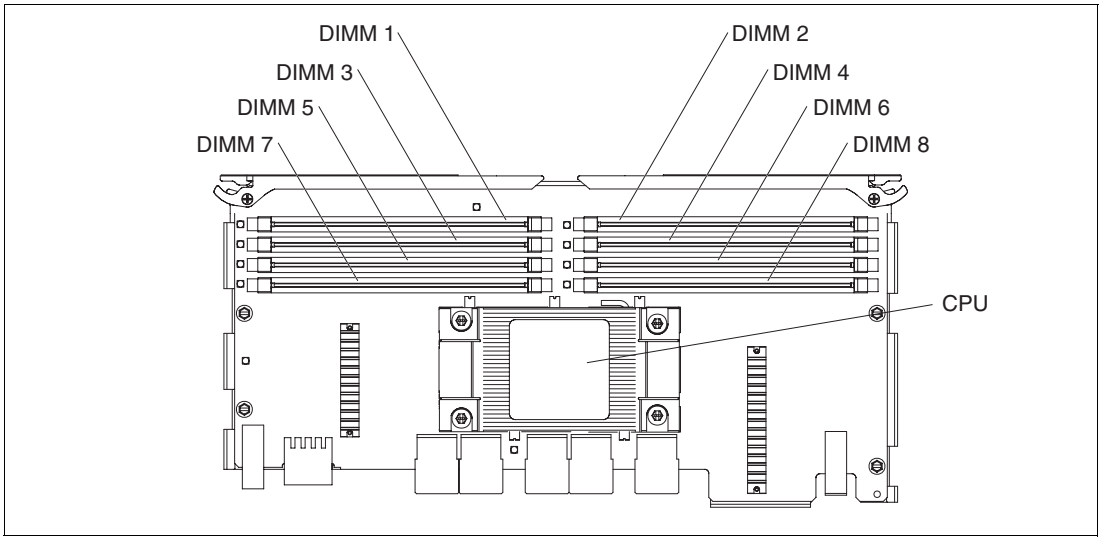

Figure 13 x3755 CPU/memory card

As discussed in “PassThru card” on page 9, a PassThru Card is required to support a configuration of either one or three processors. By using the PassThru cards, two HyperTransport buses will be connected together, eliminating the need for a CPU/memory card. Installed processors must be identical in speed and cache size.

Table 4 shows the processors that the x3755 uses.

Table 4 CPUs used by x3755 models

Model	CPUs standard	Power	Part number	Announced
8877-1RU	1x 2.0 GHz Opteron model 8212	95 W	40K1200	August 2006
8877-2RU	1x 2.2 GHz Opteron model 8214	95 W	40K1201	August 2006
8877-3RU	2x 2.4 GHz Opteron model 8216	95 W	43W7244	April 2007
8877-4RU	2x 2.6 GHz Opteron model 8218	95 W	40K1202	August 2006
8877-5RU	2x 2.8 GHz Opteron model 8220 SE	120 W	40K1203	August 2006
8877-5SU	2x 2.8 GHz Opteron model 8220	95 W	43W7246	April 2007
8877-6RU	2x 3.0 GHz Opteron model 8222 SE	120 W	43W7247	April 2007
8877-7RU	2x 3.2 GHz Opteron model 8224 SE	120 W	44R6036	August 2007

These 8000-series AMD Opteron processors are suitable for four-socket servers such as the x3755. The “2” in the second postion in the processor number indicates this is a second-generation AMD Opteron processor. The last two numbers indicate the relative performance within the 8200 series.

The processors with the label “SE” indicate higher power consumption for a specific performance level, however SE models are available before the standard models with the same processor model number. For example, the 8220 SE was available in the x3755 in August 2006, while the lower-power 8220 model became available in April 2007. The 8220 and the 8220 SE have the same performance characteristics.

All CPUs used in a server must be of the same type, speed, and L2/L3 cache size. Each CPU option includes a CPU/memory card and heatsink.

The x3755 supports the following CPU configurations:

- One CPU/memory card in slot 1 and the PassThru card in slot 2
- Two CPU/memory cards in slots 1 and 2
- Three CPU/memory cards in slot 1, 2, and 3 and the PassThru card in slot 4
- Four CPU/memory cards in slots 1, 2, 3, and 4

Note: When installing three or more CPUs, the optional Redundant Cooling Kit, part number 40K7545, is required. See “Redundant Cooling Kit” on page 24.

Figure 14 shows the internal I/O connectors and the slot numbering of the x3755.

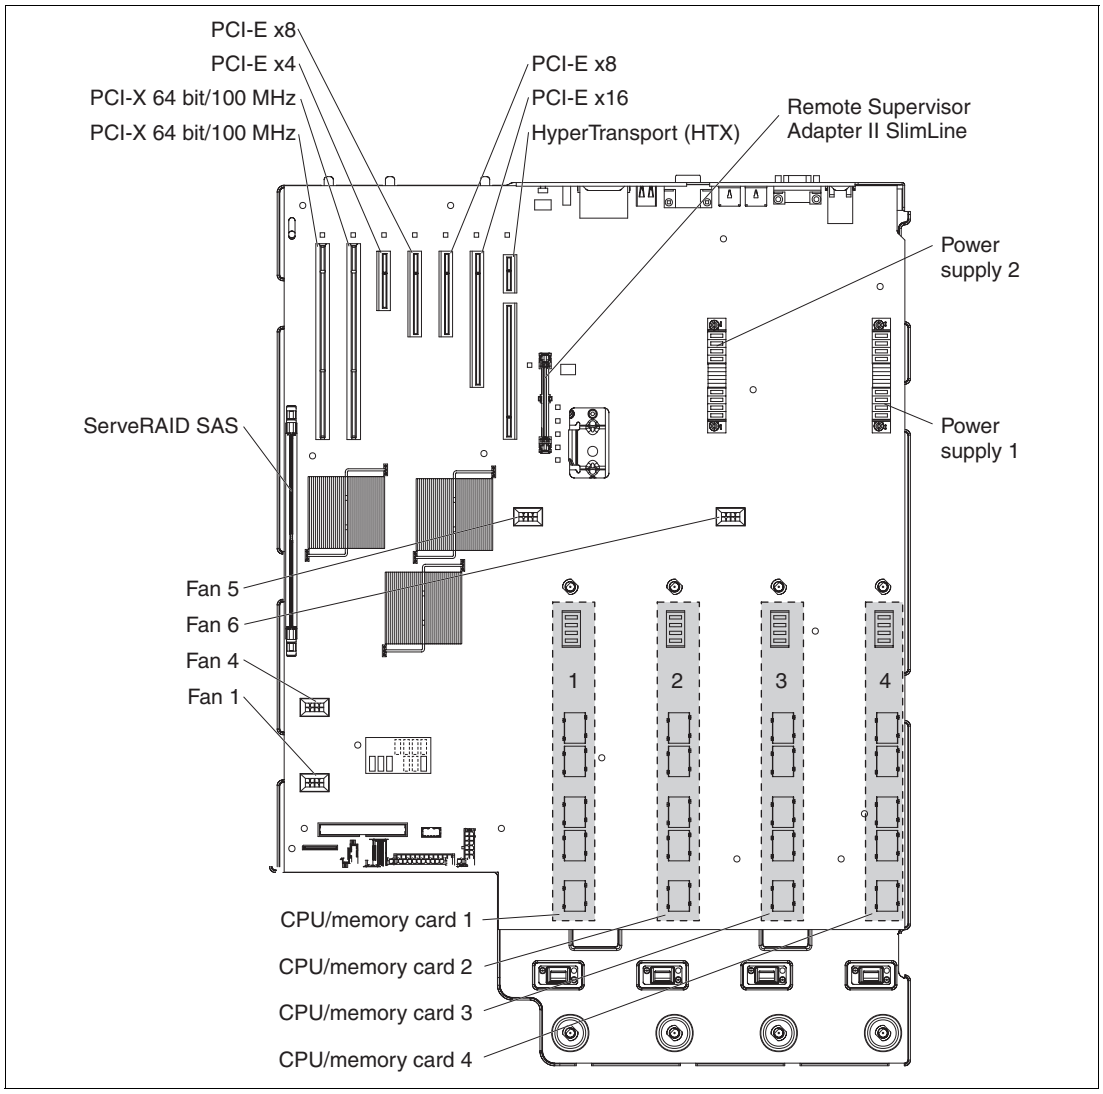

Figure 14 x3755 I/O internal connectors and CPU/memory card slot locations

The following configurations will work, but are not recommended, as they do not provide full functionality:

- One CPU/memory card in slot 1 and no PassThru card in slot 2

  As explained in “One-way configuration” on page 9, if no CPU/memory card or PassThru card is installed in slot 2, the second bridge will not be activated, and therefore the two PCI Express slots (PCI Express x16 and PCI Express x8) will not function.

- Three CPU/memory cards in slots 1, 2, and 3 and no CPU PassThru card in slot 4

  As described in “Three-way configuration” on page 10, this configuration will not take advantage of the ring topology, and therefore an extra CPU hop is inserted, which negatively affects the performance.

64-bit extensions: AMD64

The AMD AMD64 architecture extends the well-established IA32 instruction set with:

- A set of new 64-bit general purpose registers (GPR)
- 64-bit instruction pointers
- The ability to process data in 64-bit chunks
- Up to 1 TB of address space that physical memory is able to access
- 64-bit integer support and 64-bit flat virtual address space

Even though the names of these extensions suggest that the improvements are simply in memory addressability, the AMD64 are in fact fully functional 64-bit processors.

There are three distinct operation modes available in AMD64:

- 32-bit legacy mode

  The first and, in the near future, probably most widely used mode is the 32-bit legacy mode. In this mode, the AMD64 processors act just like any other IA32-compatible processor. You can install your 32-bit OS on such a system and run 32-bit applications, but you will not be able to make use of the new features such as the flat memory addressing above 4 GB or the additional General Purpose Registers (GPRs). 32-bit applications will run just as fast as they would on any current 32-bit processor.

  Most of the time, IA32 applications will run even faster since there are numerous other improvements that boost performance regardless of the maximum address size. For applications that share large amounts of data there might be performance impacts related to the NUMA-like architecture of multi-processor Opteron configurations since remote memory access might slow your application down.

- Compatibility mode

  The second mode supported by the AMD64 is compatibility mode, which is an intermediate mode of the full 64-bit mode described below. In order to run in compatibility mode, you will need to install a 64-bit operating system and 64-bit drivers. If a 64-bit OS and drivers are installed, both Opteron and Xeon® processors will be enabled to support a 64-bit operating system with both 32-bit applications or 64-bit applications.

  Compatibility mode gives you the ability to run a 64-bit operating system while still being able to run unmodified 32-bit applications. Each 32-bit application will still be limited to a maximum of 4 GB of physical memory. However, the 4-GB limit is now imposed on a per-process level, not at a system-wide level. This means that every 32-bit process on this system gets its very own 4 GB of physical memory space (assuming sufficient physical memory is installed). This is already a huge improvement compared to IA32, where the operating system kernel and the application had to share 4 GB of physical memory.

  Additionally, compatibility mode does not support the virtual 8086 mode, so real-mode legacy applications are not supported. However, 16-bit protected mode applications are supported.

- Full 64-bit mode (long mode)

  The final mode is the full 64-bit mode. AMD refer to this as *long mode*. This mode is when a 64-bit operating system and 64-bit application are use. In the full 64-bit operating mode, an application can have a virtual address space of up to 40-bits (which equates to 1 TB of addressable memory). The amount of physical memory will be determined by how many DIMM slots the server has and the maximum DIMM capacity supported and available at the time.

  Applications that run in full 64-bit mode will get access to the full physical memory range (depending on the operating system) and will also get access to the new GPRs as well as to the expanded GPRs. However it is important to understand that this mode of operation requires not only a 64-bit operating system (and of course 64-bit drivers), but also requires

a 64-bit application that has been recompiled to take full advantage of the various enhancements of the 64-bit addressing architecture.

For more information about the AMD64 architecture see:

http://www.x86-64.org/

The benefit of AMD64 64-bit computing

In the same way that 16-bit processors and 16-bit applications are no longer used in this space, it is likely that at some point in the future, 64-bit processors and applications will fully replace their 32-bit counterparts.

Processors using the AMD64 architecture are making this transition very smooth by offering 32-bit and 64-bit modes. This means that the hardware support for 64-bit will be in place before you upgrade or replace your software applications with 64-bit versions. IBM System x now has several models available with the AMD64 Opteron processors.

The question you should be asking is whether the benefit of 64-bit processing is worth the effort of upgrading or replacing your 32-bit software applications. The answer is that it depends on the application. Here are examples of applications that will benefit from 64-bit computing:

- Encryption applications

  Most encryption algorithms are based on very large integers and would benefit greatly from the use of 64-bit GPRs and ALUs. While modern high-level languages allow you to specify integers above the 2^{32} limit, in a 32-bit system, this is achieved by using two 32-bit operands, thereby causing a significant overhead while moving those operands through the CPU pipelines. A 64-bit processor will allow you to perform 64-bit integer operation with one instruction.

- Scientific applications

  Scientific applications are another example of workloads that need 64-bit data operations. Floating-point operations do not benefit from the larger integer size since floating-point registers are already 80 or 128 bits wide even in 32-bit processors.

- Software applications requiring more than 4 GB of memory

  The biggest advantage of 64-bit computing for commercial applications is the flat, potentially massive, address space.

  32-bit enterprise applications such as databases are currently implementing Page Addressing Extensions (PAEs) and Addressing Windows Extensions (AWEs) addressing schemes to access memory above the 4-GB limit imposed by 32-bit address limited processors. With EM64T and AMD64, these 32-bit addressing extension schemes support access to memory up to 128 GB in size.

  One constraint with PAE and AWE, however, is that memory above 4 GB can only be used to store data. It cannot be used to store or execute code, so these addressing schemes only make sense for applications such as databases, where large data caches are needed.

  In contrast, a 64-bit virtual address space provides for direct access to up to 2 Exabytes (EB), and even though we call these processors 64-bit, none of the current 64-bit processors actually supports a full 64 bits of physical memory addressing, simply because this is such an enormous amount of memory.

In addition, 32-bit applications might also get a performance boost from a 64-bit AMD64 system running a 64-bit operating system. When the processor runs in compatibility mode, every process has its own 4-GB memory space, not the 2-GB or 3-GB memory space each

gets on a 32-bit platform. This is already a huge improvement compared to IA32 where the OS and the application had to share those 4 GB of memory.

When the application is designed to take advantage of more memory, the availability of the additional 1 or 2 GB of physical memory can create a significant performance improvement. Not all applications take advantage of the global memory available. APIs in code need to be used to recognize the availability of more than 2 GB of memory.

Furthermore, some applications will not benefit at all from 64-bit computing and might even experience degraded performance. If an application does not require greater memory capacity or does not perform high-precision integer or floating-point operations, then 64-bit will not provide any improvement.

In fact, because 64-bit computing generally requires instructions and some data to be stored as 64-bit objects, these objects consume more physical memory than the same object in a 32-bit operating environment. The memory capacity inflation of 64-bit can only be offset by an application taking advantage of the capabilities of 64-bit (greater addressing or increased calculation performance for high-precision operations), but when an application does not make use of the 64-bit operating environment features, it often experiences the overhead without the benefit.

In this case, the overhead is increased memory consumption, leaving less physical memory for operating system buffers and caches. The resulting reduction in effective memory can decrease performance.

Software driver support in general is lacking for 64-bit operating systems compared to the 32-bit counterparts. General software drivers such as disk controllers, network adapters, or application tools might not have 64-bit code in place for x64 operating systems. Prior to moving to an x64 environment, it might be wise to ensure that all third-party vendors and software tools support drivers for the specific 64-bit operating system that you are planning to use.

64-bit memory addressing

The width of a memory address dictates how much memory the processor can address. A 32-bit processor can address up to 2^{32} bytes or 4 GB. A 64-bit processor can theoretically address up to 2^{64} bytes or 16 Exabytes (or 16,777,216 TB), although the AMD Opteron (64-bit) can directly address up to 256 TB (48-bit) and when using PAE up to 128 GB in compatibility mode.

The 64-bit extensions in the processor architectures AMD64 provide a better performance for both 32-bit and 64-bit applications on the same system. These architectures are based on 64-bit extensions to the industry-standard x86 instruction set and provide support for existing 32-bit applications.

System memory

The models of the x3755 have 1 GB or 4 GB of RAM standard, implemented using PC-5300 667 MHz registered 240-pin DDR2-SDRAM RDIMMs.

Memory is implemented in the x3755 using the CPU/memory cards (see Figure 13 on page 12). Each card has eight DIMM sockets.

The x3755 has either one CPU/memory card as standard or two, one for each CPU installed. The x3755 supports up to four CPU/memory cards total. An additional CPU/memory card is

included with each processor upgrade option. Using 4 GB DIMMs in every socket, that is a total of 32 DIMMs, the server can hold 128 GB of RAM.

The memory is two-way interleaved (meaning that memory DIMMs are installed in pairs). As shown in Figure 7 on page 8, there are four ports to memory, with each supporting up to 5.33 GBps data transfers (667 MHz x 8 bytes).

As described in "Xcelerated memory technology" on page 8, the unique IBM design allows all eight DIMMs on each CPU/memory card to operate at 667 MHz.

Key configuration rules are:

- Supported DIMM options are:
  - 1 GB (part number 41Y2759) containing two 512-MB DIMMs
  - 2 GB (part number 41Y2762) containing two 1-GB DIMMs
  - 4 GB (part number 41Y2765) containing two 2-GB DIMMs
  - 8 GB (part number 41Y2768) containing two 4-GB DIMMs
- The x3755 uses two-way interleaving memory, so DIMMs must be installed in matched pairs.
- Each CPU/memory card has eight DIMM sockets. See the *System x3755 User's Guide* for details on the required installation sequence.
- If you want to install the full 128 GB, you must remove the existing DIMMs and fully populate the x3755 with four CPU/memory cards, each with eight 4-GB DIMMs.

There are a number of advanced memory protection features implemented in the x3755 memory subsystem:

- Memory Online Spare

  A new feature, Memory Online Spare, is available on the x3755 for increased fault tolerance. The Memory Online Spare feature provides the equivalent of a hot-spare drive in a RAID array. This feature disables the failed memory DIMM if more than a predefined number of single-bit failures occur, together with the other half of its matched pair, and activates the online spare memory pair of DIMMs to replace the failed DIMM pair.

  Before you can enable this feature, you must install two additional pairs of DIMMs (or understand that useable memory will be reduced by the largest two DIMMs). The spare DIMM pairs must be the same speed, type, and size as, or larger than, the largest active DIMM pair. Memory Online Spare is enabled in the BIOS in the Memory Settings section of the Advanced Setup menu.

  Memory Online Spare is operating system independent, because all mirroring activities are handled by the hardware. After Memory Online Spare is enabled in the BIOS, any DIMM pair that fails the single bit failure threshold will be disabled and the online spare memory pair of DIMMs will replace the failed active DIMM pair.

  Important: If Memory Online Spare is enabled, the pair of DIMMs on each CPU/memory card that is designated as the online spare will not count towards the system's main memory. As Memory Online Spare is enabled for the entire system, if four CPU/memory cards are installed, a total of eight DIMMs will be designated as spare DIMMs.

- Memory scrubbing

  Memory scrubbing is an automatic daily test of all the system memory that detects and reports memory errors that might be developing before they cause a server outage.

Memory scrubbing and memory sparing work in conjunction with each other and do not require memory mirroring to be enabled to work properly.

When a bit error is detected, memory scrubbing determines whether the error is recoverable. If it is recoverable, memory sparing is enabled and the data that was stored in the damaged locations is rewritten to a new location. The error is then reported so that preventative maintenance can be performed. As long as there are enough good locations to allow the proper operation of the server, no further action is taken other than recording the error in the error logs.

If the error is not recoverable, then memory scrubbing sends an error message to the light path diagnostics, which then turns on the proper lights and LEDs to guide you to the damaged DIMM. If memory mirroring is enabled, then the mirrored copy of the data from the damaged DIMM is used until the system is powered down and the DIMM replaced.

- Chipkill memory

  Chipkill is integrated into the chipset, so it does not require special Chipkill DIMMs and is transparent to the operating system. When combining Chipkill with memory sparing, the x3755 provides very high reliability in the memory subsystem.

  When a memory chip failure occurs, Memory Online Spare transparently handles the rerouting of data around the failed component, as described above. However, if a further failure occurs, the Chipkill component in the memory controller reroutes data. The memory controller provides memory protection similar in concept to disk array striping with parity, writing the memory bits across multiple memory chips on the DIMM. The controller is able to reconstruct the missing bit from the failed chip and continue working as usual. One of these additional failures can be handled per memory port (a total of four Chipkill recoveries).

- Xcelerated Memory Technology

  As mentioned in “Xcelerated memory technology” on page 8, in addition to the above-mentioned memory protection features, the x3755 introduces a performance enhancement to the AMD architecture called Xcelerated Memory Technology.

Serial Attached SCSI

The x3755 has a disk subsystem that is comprised of an Adaptec AIC-9580W Serial Attached SCSI (SAS) controller and four internal hot-swap drive bays.

Serial Attached SCSI is the logical evolution of SCSI. SAS uses much smaller interconnects than SCSI, while offering SCSI compatibility, reliability, performance, and manageability. In addition, SAS offers longer cabling distances, smaller form factors, and greater addressability.

SAS 1.1 technology is replacing Ultra320 SCSI in SCSI and RAID controllers. Beyond the upgrades in I/O processor and memory speeds, SAS-based products will differ from SCSI-based products in the following ways:

- Higher bandwidth

  Ultra320 SCSI supports 320 MBps of bandwidth per channel. SAS 1.1 supports 3 Gbps (approximately 300 MBps) of bandwidth per port. So while the two Ultra320 SCSI channels of the ServeRAID-6M can potentially support 640 MBps of bandwidth, the onboard SAS controller with its eight ports could support up to 12 Gbps (approximately 1.2 GBps) of bandwidth. Therefore, bandwidth will be limited by PCI-X or PCI Express bus speeds.

- Greater drive support

  SCSI-based products support 14 drives for each channel. By cascading drive enclosures, SAS-based products will support up to 72 drives for every four ports. Although this is not directly relevant to the x3755 and the ServeRAID-8k because the supported SAS drives are in the four internal drive bays only, it is indicative of the advances in the technology and the future capabilities of ServeRAID SAS adapters with external storage connectivity.

The x3755 has a ServeRAID 8k-l RAID controller installed as standard. The controller is installed in a specialized socket.

The ServeRAID 8k-l has the following features:

- Support for RAID levels 0, 1, and 10
- No cache
- No battery backup

The x3755 also supports the optional ServeRAID 8k, part number 25R8064. To install the ServeRAID 8k, you must remove the ServeRAID 8k-l card. The ServeRAID 8k controller has the following features:

- Support for RAID levels 0, 1, 1E, 5, 6, 10.
- 256 MB DDR2 533 MHz memory unbuffered DIMM cache
- Battery backup
- No connectors for external SAS devices

Note: The ServeRAID slot only accepts the specialized 8k-l or the 8k RAID controllers (each implemented in the form of a memory DIMM). No industry standard DDR2 DIMMs can be used in this slot.

Graphics subsystem

The x3755 has an ATI RN50b graphics controller chip integrated on the system planar, with 8 MB of dedicated video RAM. It has the following specifications:

- 128-bit graphics engine with 8, 16, and 24 bpp mode acceleration
- 32-bpp (4G colors/true color) support
- Integrated 350 MHz RAMDAC
- DDC2B monitor communications support

In addition, the x3755 supports the NVIDIA Quadro FX 5600 Graphics Adapter, part number 43V5761, feature code 1821. This adapter has the following features:

- 1536 MB of 384-bit, GDDR3 SDRAM memory interface
- Dual 400 MHz RAMDACs supporting a primary display and a secondary analog display resolution up to 2048 x 1536 at 85 MHz per display
- Memory clock speed of 800 MHz with memory bandwidth of up to 76.8 GBps
- Dual Dual-link DVI connectors supporting digital resolutions up to 2560x1600 @ 60Hz
- 16x full-scene antialiasing (FSAA)
- NVIDIA G80GL video chip
- High-bandwidth digital content protection (HDCP) support
- 3-pin stereo mini-DIN stereo support
- Full x16 PCI-Express support
- Optimized and vendor certified for all leading workstation applications

- Hardware accelerated antialiased points and lines
- Hardware OpenGL overlay planes
- Hardware accelerated two-sided lighting
- High-performance OpenGL Hardware Accelerated Pixel Readback
- Full support for Shader Model 4.0/DirectX 10

Table 5 lists the supported resolutions, bit depth and refresh rates of the NVIDIA Quadro FX 5600.

Table 5 Supported resolutions, bit depth and refresh rates of the NVIDIA Quadro FX 5600

Resoltuion	Color bit depth	Refresh rates
320 x 200	8, 16, 32	60, 70, 75
320 x 240	8, 16, 32	60, 70, 75
400 x 300	8, 16, 32	60, 70, 75
480 x 360	8, 16, 32	60, 70, 75
512 x 384	8, 16, 32	60, 70, 75
640 x 400	8, 16, 32	60, 70, 72, 75, 85, 100, 120
640 x 480	8, 16, 32	60, 70, 75, 85, 100, 120
800 x 600	8, 16, 32	60, 70, 72, 75, 85, 100, 120
1024 x 768	8, 16, 32	60, 70, 72, 75, 85, 100, 120
1152 x 864	8, 16, 32	60, 70, 72, 75, 85, 100, 120
1280 x 720	8, 16, 32	60, 70, 72, 75, 85, 100, 120
1280 x 960	8, 16, 32	60, 70, 75, 85, 100, 120
1280 x 1024	8, 16, 32	60, 70, 72, 75, 85, 100, 120
1600 x 900	8, 16, 32	60, 70, 75, 85, 100
1600 x 1200	8, 16, 32	60, 70, 72, 75, 85, 100
1792 x 1344	8, 16, 32	60, 75
1856 x 1392	8, 16, 32	60, 75
1920 x 1080	8, 16, 32	60, 70, 72, 75, 85, 100
1920 x 1200	8, 16, 32	60, 70, 72, 75, 85,
1920 x 1440	8, 16, 32	60, 70, 72, 75, 85
2048 x 1536	8, 16, 32	60, 70, 72, 75, 85

Table 6 shows the maximum resolutions supported.

Table 6 Maximum resolutions supported

	Maximum analog resolution	Maximum digital resolution
Single head	2048 x 1536	1600 x 1200 (60 Hz)
Dual-head	2048 x 1536 each	1600 x 1200 each (60 Hz)

Only one NVIDIA Quadro FX5600 Graphics Adapter is supported and it must be installed in slot 1 of the x3755.

PCI subsystem

The x3755 has six PCI slots as follows:

- One full-length PCI Express x16 slot
- Two full-length PCI Express x8 slots
- One full-length PCI Express x4 slot
- Two full-length 64-bit, PCI-X 1.0, 100/133 MHz slots
- One half-length HTX slot

Important: PCI slot 1 and 2 (PCI Express x16 and PCI Express x8) are only enabled if either a CPU/memory card or a PassThru card is installed in CPU slot 2. PCI slot 7 (HTX) is only enabled if at least three CPU/memory card are installed. See “PassThru card” on page 9.

Figure 4 on page 5 shows the six PCI slots found in the x3755.

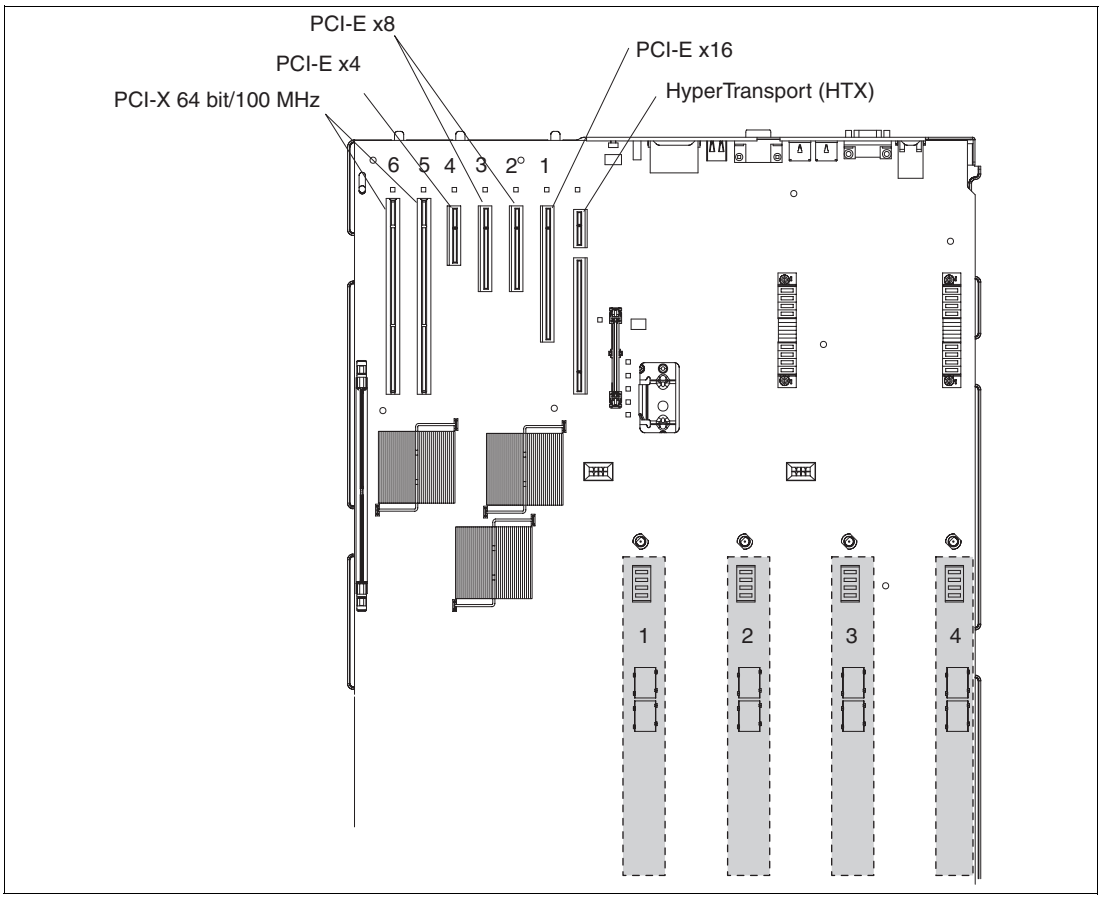

Figure 15 PCI slots

Note: The optional Redundant Cooling Kit 40K7545 is needed if the client wants redundant fan cooling for the PCI slots.

Details about these slots are:

- The PCI Express slots have the following characteristics:
  - PCI Express 2.0
  - 266 MHz, each supporting lower-speed adapters
  - 64-bit, each supporting 32-bit adapters as well
  - 3.3 V
  - Separate bus from the other slots and devices, meaning that the speed of the adapter does not affect the other adapters
- The PCI-X slots share the same bus they will run at the speeds indicated by Table 7.

Table 7 Bus speed of the 2 PCI-X slots

Slot 5	Slot 6	Bus Speed
PCI-X 133 MHz card installed	No card installed	133 MHz
No card installed	PCI-X 133 MHz card installed	100 MHz
PCI-X 133 MHz card installed	PCI-X 133 MHz card installed	100 MHz

- The half length HTX slot that can accommodate cards such as the QLogic InfiniPath HTX InfiniBand Adapter. These cards will extend the high-speed HyperTransport CPU-to-CPU interconnect across connected systems. As shown in Figure 4 on page 5, this slot is connected to CPU 3 without a bridge. Therefore devices connected through the HTX slot can transfer data to and from main system memory without incurring any penalty normally associated with passing data through a bridge.

Note: To enable the HTX slot, you must have either three CPU/memory cards with a PassThru card or four CPU/memory cards installed.

One additional PCI slot is reserved for the ServeRAID-8k adapter.

The PCI subsystem also supplies these I/O devices:

- Adaptec AIC-9580W Serial-attached SCSI (SAS) controller
- Broadcom dual port 5708C 10/100/1000 Ethernet
- ATI ES1000 video controller (with 16 MB video memory)
- Three USB ports (two on the front panel, one on the rear)
- EIDE interface for the optical media drive
- shared RSA/System Serial port
- Remote Supervisor Adapter II SlimLine adapter (in a dedicated socket on the I/O board)

Note: There is no parallel port on the x3755. For parallel port connections, use the Lenovo USB Serial/Parallel Adapter, part number 22P5298. In addition, there is no diskette drive. If a diskette drive is required, use the IBM USB Portable Diskette Drive, part number 40K1692.

Redundant Cooling Kit

The x3755 has three fans as standard. Fan 1 is installed in the I/O zone to cool the PCI slots, and fans 2 and 5 are installed to provide redundant cooling for CPU slots 1 and 2. Figure 16 shows the fan zones and the fan locations.

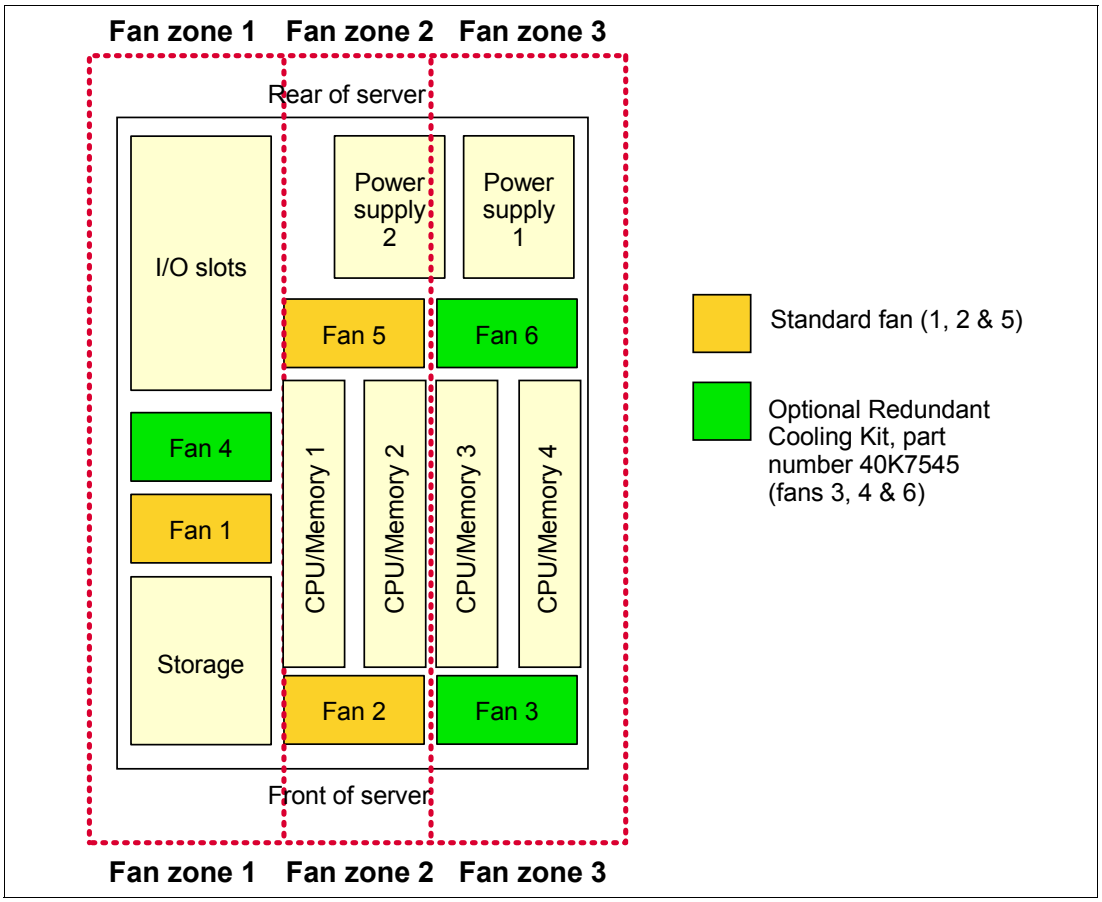

Figure 16 x3755 fan zones and fan location

The optional Redundant Cooling Kit, part number 40K7545, provides an additional three fans. One fan is installed in position 4 in the I/O zone to provide redundant cooling for the PCI slots, and the other two fans are in placed positions 3 and 6 and are used to provide the redundant fans for CPU slots 3 and 4. Therefore the Redundant Cooling Kit is required in the following configurations:

- A CPU/memory card is installed in slot 3 or 4.
- The client requires redundancy in the I/O slot section.

Each fan zone — once the kit is installed — will provide cooling redundancy for the individual sections, thus enabling individual fan replacement without powering down the server. In the event of a fan failure, the redundant fan will speed up to continue to provide adequate cooling until the fan can be hot-swapped by the IT administrator. In general, failed fans should be replaced within 48 hours following failure.

Redundancy

The x3755 has the following redundancy features to maintain high availability:

- There are three hot-swap multi-speed fans standard. One fan is installed in the I/O slot section of the system. Two fans are installed in the first CPU section.

  The optional Redundant Fan Kit provides another three fans. See “Redundant Cooling Kit” on page 24 for details when this kit is required.
- The two Gigabit Ethernet ports can be configured as a team to form a redundant pair.
- The memory subsystem has a number of redundancy features, including Chipkill and Memory Online Spare, as described in “System memory” on page 17.
- Support is available for RAID disk arrays, both with the ServeRAID-8k (for internal RAID arrays) and the MegaRAID 8480 (for SAS external arrays). The x3755 has four internal hot-swap disk drive bays.
- An optional second 1500 W hot-swap power supply is available and connected using a separate power cord. At 220 V, this second power supply is redundant in all configurations.

Important: At 110 V and with a heavily configured server, two power supplies may be needed. The second power supply will not be redundant.

Light path diagnostics

To limit the need to slide the server out of the rack to diagnose problems, a light path diagnostics panel is located at the front of the x3755. This panel slides out from the front of the server so that the client can view all light path diagnostics monitored server subsystems. In the event that maintenance is required, the client can slide the server out from the rack and, using the LEDs, find the failed or failing component.

Light path diagnostics can monitor and report on the health of CPUs, main memory, hard disk drives, PCI-X and PCI adapters, fans, power supplies, and the internal system temperature.

Baseboard Management service processor

The Baseboard Management Controller (BMC) is a small, independent micro-controller used to perform low-level system monitoring and control functions, as well as remote IPMI interface functions. It uses multiple I2C bus connections to communicate out-of-band with other onboard devices. The BMC provides environmental monitoring for the server. If environmental conditions exceed thresholds or if system components fail, the BMC lights the light path diagnostic LEDs to help you diagnose the problem and also records the error in the BMC system event log.

BMC functions are as follows:

- Initial system check at A/C on

  The BMC monitors critical I2C devices in standby power mode to determine whether the system configuration is safe for power-on.
- BMC event log maintenance

  The BMC maintains and updates an IPMI-specified event log in non-volatile storage. Critical system information is recorded and made available for external viewing.

- System power state tracking

  The BMC monitors the system power state and logs transitions into the system event log.

- System initialization

  The BMC has I2C access to certain system components that may require initialization before power-up.

- System software state tracking

  The BMC monitors the system and reports when the BIOS and POST phases are complete and the operating system has booted.

- System event monitoring

  During runtime, the BMC continually monitors critical system items such as fans, power supplies, temperatures, and voltages. The system status is logged and reported to the service processor, if present.

- System fan speed control

  The BMC monitors system temperatures and adjusts fan speed accordingly.

The BMC also provides the following remote server management capabilities through the OSA SMBridge management utility program:

- Command-line interface (IPMI Shell)
- Serial over LAN (SOL)

For more information about how to enable and configure these management utilities, see the *x3755 User's Guide.*

Remote Supervisor Adapter II SlimLine

The x3755 can be upgraded with the Remote Supervisor Adapter II SlimLine service processor, part number 39Y9566, to enable full RSA II support. This adapter, which installs in a dedicated slot, provides the similar functionality as the Remote Supervisor Adapter II PCI option available for other System x™ servers.

Note: The previous RSA II SlimLine, part number 73P9341, is not supported.

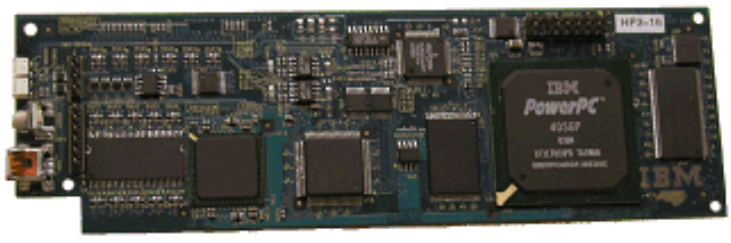

Figure 17 Remote Supervisor Adapter II SlimLine daughter card

Key features of the Remote Supervisor Adapter II SlimLine include:

- IBM ASIC with integrated PowerPC® 405 core executing at 200 MHz
- 16 MB SDRAM and 4 MB flash ROM

- System-independent graphical console redirection
  - Built-in video compression hardware eliminates drivers.
  - Graphics response, up to five times faster than with the RSA I, makes monitoring and control more efficient.
  - System-independent installation eliminates the need to install service processor drivers, helps save IT staff time, and reduces installation complexity.
- Remote diskette and CD-ROM drive support
  - Enables remote booting and software loading of the server for application/operating system installation and updates.
  - Performs configuration remotely; helps save IT time and money by reducing on-site presence and server downtime.
- Scriptable command-line interface and text-based serial console redirect
  - Command-line interface supports program control of server management functions using scripts.
  - Serial text redirect provides access to text-mode BIOS and text-based system consoles such as Linux, NetWare, and Windows EMS (Emergency Management Services).
  - Program control of text-based console using scripts.
  - PPP support.
- User authentication and authority features
  - User IDs, passwords, and login permission attributes can be stored in an LDAP server.
  - Enhanced user authority levels set the access rights for users to match job responsibilities for managing your IBM System x servers.
  - Secure Sockets Layer (SSL) encrypts the data transmitted between the LDAP server and the Remote Supervisor AdapterII SlimLine.
- Investment protection
  - Integrates with IBM Director and Director Agent.

PowerExecutive

IBM PowerExecutive™ is a combination of hardware and software that allows direct power monitoring through IBM Director and helps automate the management of power consumption. The x3755 supports PowerExecutive. This tool helps clients monitor power consumption to allow better utilization of available power resources.

This application software enables clients to trend actual power consumption and corresponding thermal loading of servers running in their environment with their applications.

IBM PowerExecutive enables clients to monitor actual power draw and thermal loading information. This helps clients with:

- More efficient planning of new data center construction or modification
- Proper power input sizing based on physical systems
- Justification of incremental hardware purchases based on available input power capacity
- Better utilization of existing resources

See the following for more information and to download the tool:

http://www.ibm.com/systems/management/director/extensions/powerexec.html

Supported operating systems

The following operating systems are planned to be supported either at general availability or soon after:

- Microsoft® Windows:
  - Windows 2000 Server SP4
  - Windows 2000 Advanced Server SP4
  - Windows Server® 2003 R2, Standard x64 Edition
  - Windows Server 2003 R2, Enterprise Edition
  - Windows Server 2003 R2, Enterprise x64 Edition
- Red Hat Enterprise Linux (RHEL)
  - Red Hat Enterprise Linux 4 U4 AS, ES, WS
  - Red Hat Enterprise Linux 4 U4 AS, ES, WS for x64
- SUSE Linux Enterprise Server (SLES)
  - SUSE Linux Enterprise Server 10 (Memory limitation of 48 GB)
  - SUSE Linux Enterprise Server 10 64-bit
  - SUSE Linux Enterprise Server 9 SP3
  - SUSE Linux Enterprise Server 9 SP3 for x64

Note: At the time of writing only Linux drivers were available for the HTX slot.

Other operating systems:

- VMware ESX Server 3.0 Update 1
- Novell NetWare 6.5

Specific service packs or updates may be required, and some operating systems may be limited as to the number of supported processors.

For the latest operating system support information see the ServerProven® operating system support matrix:

http://www.ibm.com/servers/eserver/serverproven/compat/us/nos/matrix.shtml

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

David Watts is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces IBM Redbooks on hardware and software topics related to IBM System x servers and associated client platforms. He has authored more than 50 IBM Redbooks, IBM Redpapers and IBM Technotes. He holds a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for more than 15 years. He is an IBM Certified IT Specialist.

Robert Moon is the Team Leader at IBM System x Techline in Greenock, Scotland. He is coauthor of the *IBM System x3850 Solution Assurance Product Review Guide*. He has more than ten years of experience in System x servers, including presales and post-sales technical support. He has been involved with Solution Assurance for over eight years, and he is currently the EMEA coordinator for the ServerProven Opportunity Request for Evaluation (SPORE) program. His areas of expertise include Novell NetWare and System x hardware. He is an IT Advisory Systems Specialist, chartered professional of the British Computer

Society (MBCS CITP), Certified Novell NetWare Administrator (CNA), IBM eServer™ Certified specialist for xSeries®, holds a degree in Business Administration and Engineering from the University of Eindhoven, and has worked for IBM for more than 10 years.

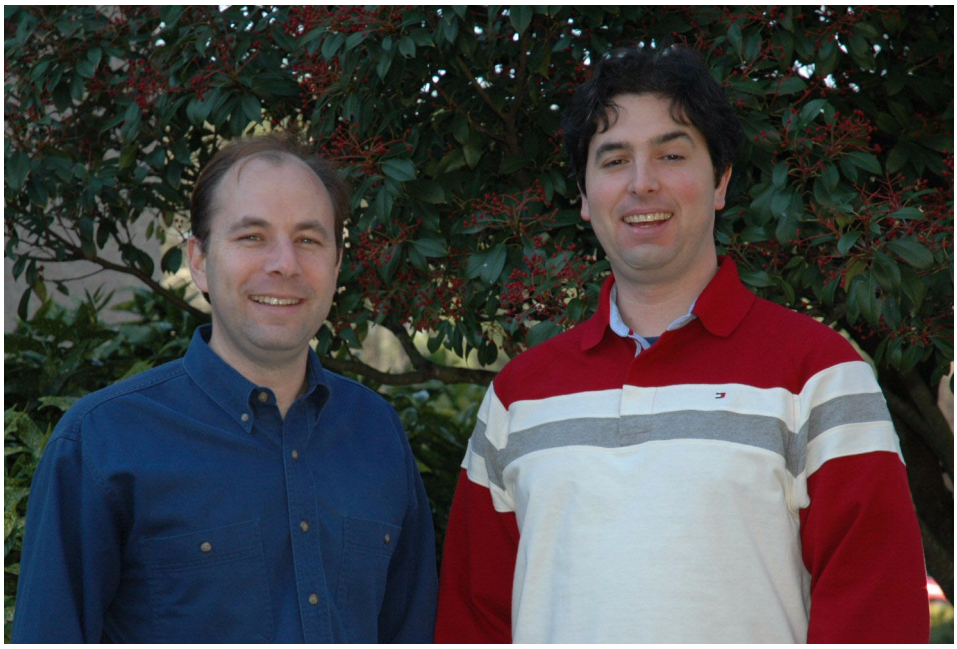

Figure 18 The team (left–right): David and Robert

Thanks to the following people for their contributions to this project:

Jay Bretzmann, Worldwide Marketing Manager, High Performance System x, Raleigh
Barney Hallman, High-End System x Server Development, Raleigh
Kyle Hampton, Worldwide Marketing Strategist, Raleigh
Dan Hurlimann, Senior Engineer, System x3755 Development, Raleigh
Beth McElroy, Worldwide Product Marketing Manager, Raleigh
Jim Marschausen, Worldwise SPORE Program Manager
Fraser Lawrence, System x Techline Team Leader
Julie Czubik, Technical Editor, International Technical Support Organization, Poughkeepsie

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document created or updated on September 5, 2007.

Send us your comments in one of the following ways:
- Use the online **Contact us** review form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo) ®	IBM®	ServerProven®
eServer™	PowerExecutive™	ServeRAID™
xSeries®	PowerPC®	System x™
Chipkill™	Redbooks®	Xcelerated Memory Technology™

The following terms are trademarks of other companies:

Microsoft, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.