



David Watts

# IBM *@*server xSeries 365 Technical Introduction

## Introduction

The four-way Intel Xeon MP-based x365 is the follow-on to the popular x360 and offers second-generation Enterprise X-Architecture technology in a rack-dense 3U form factor. The x365 is targeted at database, e-mail, e-commerce and ERP applications.

This paper introduces the new xSeries 365 and describes the technical aspects of the server. Topics in this paper are:

- ▶ “Overview of the x365” on page 1
- ▶ “Current models” on page 2
- ▶ “IBM XA-32 second-generation chipset” on page 3
- ▶ “Processors” on page 5
- ▶ “System memory” on page 8
- ▶ “PCI subsystem” on page 13
- ▶ “RXE-100 Remote Expansion Enclosure” on page 18
- ▶ “Redundancy” on page 22
- ▶ “Light path diagnostics” on page 23
- ▶ “Operating system support” on page 23

This document has been updated to correct some technical inaccuracies. Changes are marked with change bars.

## Overview of the x365

The following are the key features of the x365:

- ▶ XA-32 second-generation chipset.
- ▶ Models with one or two Intel Xeon MP processors, up to 2.8 GHz and 2 MB L3 cache, upgradable to four-way.

- ▶ 1 GB or 2 GB memory standard, expandable to 32 GB (16 DIMM slots), using high-performance ECC DDR DIMMs.
- ▶ Active Memory with Memory ProteXion, memory mirroring and ChipKill.
- ▶ Four 133 MHz and one 100 MHz Active PCI-X slots.
- ▶ Integrated LSI 1030 dual-channel Ultra320 SCSI controller with integrated RAID-1 support. Addition support for internal RAID arrays using an optional ServeRAID adapter.
- ▶ Six internal hot-swap drive bays for up to 876 GB of internal storage.
- ▶ Integrated Dual Broadcom 5704 PCI-X Gigabit Ethernet.
- ▶ Support for an RXE-100 Remote Expansion Enclosure for up to 12 additional PCI-X slots. The RXE-100 can be used exclusively by one x365 or can be shared between two x365s (six slots to each server).
- ▶ 3U rack-optimized tool-free chassis.
- ▶ Remote Supervisor Adapter II standard.
- ▶ Supports the IBM Integrated xSeries Adapter for iSeries (IXA) for a direct high speed link to an iSeries server.
- ▶ One-year warranty (for machine type 8861) or 3-year warranty (for machine type 8862). Standard warranty is on-site, 9 hours per day, 5 days per week, next business day response.

Figure 1 shows the x365 and major components on the front of the unit.




Figure 1 Front panel of the x365

## Current models

The x365 models, shown in Table 1, were announced in December 2003.

Table 1 x365 models

Model	Standard/max CPUs	L2 cache	L3 cache	Standard/max memory	Std/max power
8861-1RX	1x Xeon MP 2.0 GHz / 4	512 KB	1 MB	1 GB (2x 512 MB) / 32	1x 950 W / 2 max
8862-1RX	1x Xeon MP 2.0 GHz / 4	512 KB	1 MB	2 GB (4x 512 MB) / 32	1x 950 W / 2 max
8862-2RX	2x Xeon MP 2.5 GHz / 4	512 KB	1 MB	2 GB (4x 512 MB) / 32	2x 950 W / 2 max
8862-3RX	2x Xeon MP 2.8 GHz / 4	512 KB	2 MB	2 GB (4x 512 MB) / 32	2x 950 W / 2 max

The x365 supports one, two, three, or four processors.

All models support a maximum of 32 GB using 2 GB DIMMs in 16 sockets. However, models other than 8862-3RX require the replacement of the 8-socket memory board to support 16 DIMMs.

The x365 has five internal PCI-X slots and one internal PCI slot. The attachment of a single RXE-100 Remote Expansion Enclosure is supported, providing an additional six or 12 PCI-X slots. Connectivity is via a single RXE expansion connection as shown in Figure 2. Furthermore, the PCI-X slots in a single RXE-100 can be shared between two x365 servers. See “RXE-100 Remote Expansion Enclosure” on page 18 for more details.

The RXE-100 has six PCI-X slots standard, upgradable to 12 PCI-X slots, giving the customer up to a total of 11 PCI-X or 17 PCI-X slots, respectively.




Figure 2 x365 configurations

## IBM XA-32 second-generation chipset

The x365 uses the second generation of the IBM XA-32 chipset, which also powers the x445 and x455 high-performance servers. The architecture consists of the following components:

- ▶ One to four Xeon MP processors
- ▶ One Cyclone 3.0 Memory Controller
- ▶ Two Winnipeg 4.0 PCI Bridges

Figure 3 shows the block diagram of the x365.

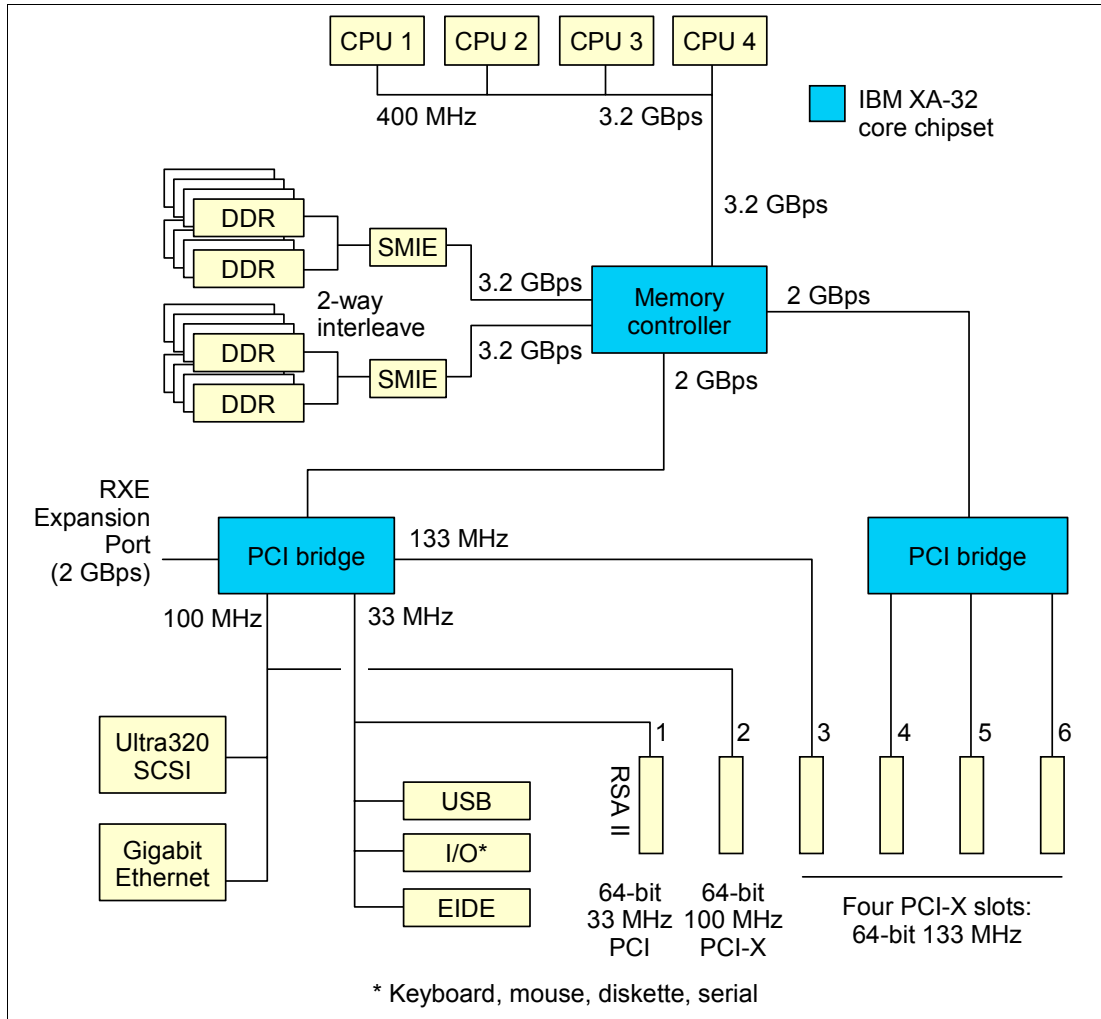


Figure 3 xSeries 365 system block diagram

One PCI bridge supplies three of the four 64-bit 133 MHz PCI-X slots on three independent PCI-X buses. The other PCI bridge supplies the other three PCI-X slots (all 64-bit, but the bus speeds vary), all the onboard PCI devices, plus the interface to the optional RXE-100 expansion enclosure.

This design is similar to the x445 and x455 with the following major differences:

- ▶ No SMP scalability ports
- ▶ No Cache and Scalability Controller
- ▶ No Xcel4 Server Accelerator Cache
- ▶ Different PCI-X slot configuration
- ▶ RSA II occupies one of the six PCI-X slots, not a separate dedicated PCI slot

The rear panel of the x365 is shown in Figure 4.




Figure 4 Rear view of the x365

## Processors

The x365 uses Intel Two Xeon MP “Gallatin” processors. The x365 use different processors as described in Table 2; however, all models support one, two, three, or four processors.

Table 2 Processors used in the x365 models

Model	Processor	L2 cache	L3 cache	Installed	Upgrade option
8861-1RX	Xeon MP 2.0 GHz	512 KB	1 MB	1	73P7073
8862-1RX	Xeon MP 2.0 GHz	512 KB	1 MB	1	73P7073
8862-2RX	Xeon MP 2.5 GHz	512 KB	1 MB	2	73P7074
8862-3RX	Xeon MP 2.8 GHz	512 KB	2 MB	2	73P7075

The upgrade options include one processor and one VRM. The VRM must be installed if the option is to be installed as CPU 3 or CPU 4. The VRMs for CPUs 1 and 2 are integrated onto the processor board.

All CPUs used in a server must be the same type, speed, and L2/L3 cache size.

The processors are easily accessible from the front of the server on a sliding tray.




Figure 5 x365 processor tray

The VRMs for processors 1 and 2 are integrated on the microprocessor board; the VRMs for processors 3 and 4 come with the processor options.

See “Processors” on page 5 for a further discussion about what you should consider before implementing an x365 solution.

The Xeon Processor MP (code named “Gallatin”) uses the ZIF socket design. This smaller form factor means that the x365 can have up to four processors in a 3U chassis.

The Xeon MP processor has three levels of cache, all of which are on the processor die:

- ▶ Level 3 cache is equivalent to L2 cache on the Pentium III Xeon. The x365 processors contain either 1 MB or 2 MB L3 cache.
- ▶ Level 2 cache is equivalent to L1 cache on the Pentium III Xeon and is 512 KB in size. The L2 cache implements the Advanced Transfer Cache technology, which means L2-to-processor transfers occur across a 256-bit bus in only one clock cycle.
- ▶ A new Level 1 execution trace cache, 12 KB in size, is “closest” to the processor and is used to store micro-operations (that is, decoded executable machine instructions); it serves those to the processor at rated speed. This additional level of cache saves decode time on cache hits. There is an additional 8 KB data cache for data related to those instructions, which is clocked at twice the speed of the processor’s clock rate.

For more information, see the Intel Technology Overview white paper, available from:

<http://www.intel.com/eBusiness/pdf/prod/server/xeon/ds020903.pdf>

Intel has also introduced a number of features associated with its NetBurst micro-architecture. These are available in the x365, including:

- ▶ 400 MHz frontside bus

The Pentium III Xeon processor in older servers had a 100 MHz frontside bus that equates a burst throughput of 800 MBps. With protocols such as TCP/IP, this has been shown to be a bottleneck in high-throughput situations.

The Xeon Processor MP improves on this by using two 100 MHz clocks, out of phase with each other by 90 degrees, and using both edges of each clock to transmit data. This is shown in Figure 6 on page 7.




Figure 6 Quad-pumped frontside bus

This increases the performance of the frontside bus without the difficulty of high-speed clock signal integrity issues. Since the bus is eight bytes wide, the end result is an effective burst throughput of 3.2 GBps, which can have a substantial impact, especially on TCP/IP-based LAN traffic.

#### ► Hyper-Threading

Hyper-Threading technology enables a single physical processor to execute two separate code streams (threads) concurrently. To the operating system, a processor with Hyper-Threading appears as two *logical* processors, each of which has its own architectural state, that is, its own data, segment and control registers, and its own advanced programmable interrupt controller (APIC).

Each logical processor can be individually halted, interrupted, or directed to execute a specified thread, independently from the other logical processor on the chip. Unlike a traditional two-way SMP configuration that uses two separate physical processors, the logical processors share the execution resources of the processor core, which includes the execution engine, the caches, the system bus interface, and the firmware.

**Note:** Hyper-Threading is enabled by default on the x365, but can be easily disabled in the BIOS.

Hyper-Threading technology is designed to improve server performance by exploiting the multi-threading capability of operating systems, such as Windows 2003 and Linux, and server applications, in such a way as to increase the use of the on-chip execution resources available on these processors.

Fewer or slower processors usually achieve the best gains from Hyper-Threading because there is a greater likelihood that the software can spawn sufficient numbers of threads to keep both paths busy. The following performance gains are likely:

- Two physical processors: 15–25 percent performance gain
- Four physical processors: 1–13 percent gain
- Eight physical processors: 0–5 percent gain

Tests have shown that software often limits SMP scalability, but customers should expect improved results as software matures. Best-case applications today are:

- Databases
- Java
- Web servers
- E-mail

Note that Microsoft licensing of the Windows Server 2000 operating systems is by number of processors (four-way for Server, eight-way for Advanced Server, 32-way for Datacenter Server). Therefore, the appearance of twice as many logical processors can potentially affect the installation of the operating system. See “Operating system support” on page 23 for details.

Windows 2003 Server operating systems understand the concept of physical processors versus logical processors. For more information about Hyper-Threading, see the following URLs:

<http://www.intel.com/technology/hyperthread/>

<http://www.microsoft.com/windows2000/server/evaluation/performance/reports/hyperthread.asp>

► **Advanced Dynamic Execution**

The Pentium III Xeon in older servers had a ten-stage pipeline. However, the large number of transistors in each pipeline stage meant that the processor was limited to speeds under 1 GHz due to latency in the pipeline.

The Xeon Processor MP has a 20-stage pipeline, which can hold up to 126 concurrent instructions in flight and up to 48 reads and 24 writes active in the pipeline. The lower complexity of each stage also means that future clock speed increases are possible.

It is important to note, however, that the longer pipeline means that it now takes more clock cycles to execute the same instruction as compared to the Pentium III Xeon.

For more information about the features of the Xeon Processor MP, go to:

<http://www.intel.com/design/Xeon/xeonmp/prodbref/>

## System memory

The x365 has one or two GB standard, implemented as two or four 512 MB PC266 (also known as PC2100) ECC DDR DIMMs.

As standard, the x365 supports either eight or 16 DIMMs depending on the model, as listed in Table 3, although models supporting eight DIMMs can be upgraded. Features such as Memory ProteXion and memory mirroring are only available when the server has 16 DIMM sockets.

The x365 comes standard with eight or 16 DIMM sockets, depending on the model, as listed in Table 3.

*Table 3 Memory DIMMs in x365 models*

Model	Standard DIMM sockets	Standard RAM	Upgrade to 16 sockets
8861-1RX	8	1 GB (2x 512 MB)	Use Memory Upgrade Card, 73P9710
8862-1RX	8	2 GB (4x 512 MB)	Use Memory Upgrade Card, 73P9710
8862-2RX	8	2 GB (4x 512 MB)	Use Memory Upgrade Card, 73P9710
8862-3RX	16	2 GB (4x 512 MB)	Not needed

With 16 sockets and using 2 GB DIMMs, the x365 can have up to 32 GB RAM.



**Important:** The ramifications of the number of sockets goes beyond simply the amount of RAM that can be installed without also replacing the memory card. Active Memory features Memory ProteXion and memory mirroring are only available when the server has 16 DIMM sockets.

The memory is two-way interleaved (meaning that memory DIMMs are installed in pairs) with two ports to the memory controller supporting up to 6.4 GBps data transfers.

There are a number of advanced features implemented in the x365 memory subsystem, collectively known as *Active Memory*:

► Memory ProteXion

Memory ProteXion, also known as “redundant bit steering,” is the technology behind using redundant bits in a data packet to provide backup in the event of a DIMM failure.

Currently, other industry-standard servers use 8 bits of the 72-bit data packets for ECC functions and the remaining 64 bits for data. However, the x365 needs only 6 bits to perform the same ECC functions, thus leaving 2 bits free. In the event that a chip failure on the DIMM is detected by memory scrubbing, the memory controller can re-route data around that failed chip through the spare bits (similar to the hot-spare drive of a RAID array). It can do this automatically without issuing a Predictive Failure Analysis (PFA) or light path diagnostics alert to the administrator. After the second DIMM failure, PFA and light path diagnostics alerts would occur on that DIMM as normal.

**Important:** Memory ProteXion is only available when the 16-socket memory board is installed (either standard, or via the Memory Upgrade Card, 73P9710).

► 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 ProteXion 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 if the error is recoverable or not. If it is recoverable, Memory ProteXion 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 in the damaged DIMM is used until the system is powered down and the DIMM replaced.

► Memory mirroring

Memory mirroring is roughly equivalent to RAID-1 in disk arrays, in that memory is divided in two ports and one port is mirrored to the other half (see Figure 8 on page 11). If 8 GB is installed, then the operating system sees 4 GB once memory mirroring is enabled (it is disabled in the BIOS by default). Since all mirroring activities are handled by the hardware, memory mirroring is operating system independent.

Certain restrictions exist with respect to placement and size of memory DIMMs when memory mirroring is enabled. See “Memory mirroring” on page 12 for details.

**Important:** Memory mirroring is only available when the 16-socket memory board is installed (either standard, or via the Memory Upgrade Card, 73P9710).

► Chipkill memory

Chipkill is integrated into the XA-32 second-generation chipset and does not require special Chipkill DIMMs. Chipkill corrects multiple single-bit errors to keep a DIMM from failing. When combining Chipkill with Memory ProteXion and Active Memory, the x365 provides very high reliability in the memory subsystem. Chipkill memory is approximately 100 times more effective than ECC technology, providing correction for up to four bits per DIMM (eight bits per memory controller), whether on a single chip or multiple chips.

If a memory chip error does occur, Chipkill is designed to automatically take the inoperative memory chip offline while the server keeps running. 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.

Chipkill support is provided in the memory controller and implemented using standard ECC DIMMs, so it is transparent to the OS.

In addition, to maintain the highest levels of system availability, if a memory error is detected during POST or memory configuration, the server can automatically disable the failing memory bank and continue operating with reduced memory capacity. You can manually re-enable the memory bank after the problem is corrected via the Setup menu in the BIOS.


Memory mirroring, Chipkill, and Memory ProteXion provide multiple levels of redundancy to the memory subsystem. Combining Chipkill with Memory ProteXion enables up to two memory chip failures per memory port on the x365, for a total of four failures sustained.

1. The first failure detected by the Chipkill algorithm on each port does not generate a light path diagnostics error, since Memory ProteXion recovers from the problem automatically.
2. Each memory port could then sustain a second chip failure without shutting down.
3. Provided that memory mirroring is enabled, the third chip failure on that port would send the alert and take the DIMM offline, but keep the system running out of the redundant memory bank.

The DIMM sockets are houses in the memory cassette and accessible through two clear doors. You do not need to remove the memory cassette to install DIMMs.



The memory cassette partially removed from the server chassis (eight-DIMM board). The clear plastic doors open for easy access to DIMMs.



The eight-DIMM board, from above

Figure 7 x365 memory cassette

The DIMM sockets are divided into two ports. When memory mirroring is enabled (16-DIMM socket board only), memory is mirrored from one port to the other.

The DIMM sockets are also divided up into banks. A bank is the pair of DIMMs required for two-way interleaving. The banks are shown in Figure 8.

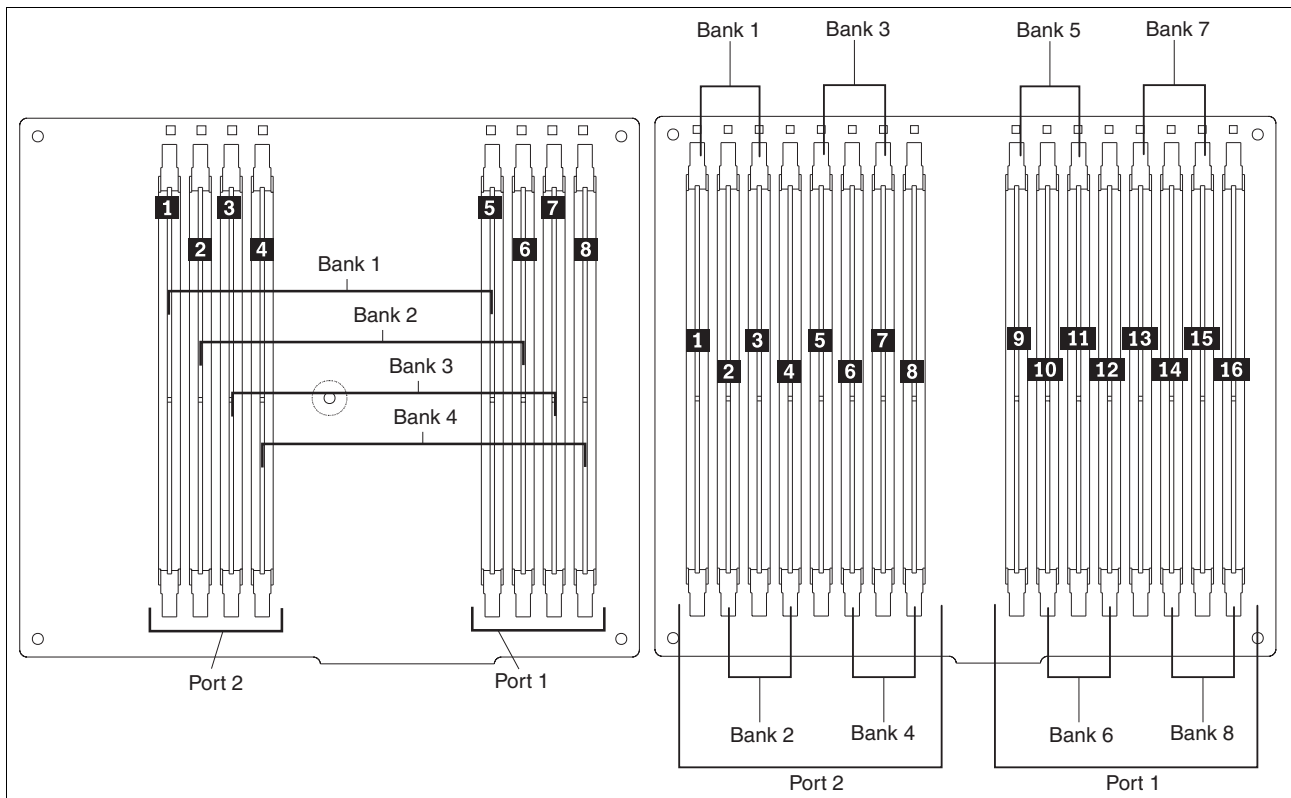


Figure 8 Eight-DIMM board and 16-DIMM board showing bank and port arrangements

Key configuration rules are:

- ▶ Because the x365 uses DDR (two-way interleaving) memory, DIMMs must be installed in banks (two DIMMs). Supported DIMMs are:
  - 512 MB DIMMs (part number 33L5038)
  - 1 GB DIMMs (part number 73P2031)
  - 2 GB DIMM (part number 73P2030) (Note: 33L5040 is not supported)
- ▶ The above DIMMs can all be used at the same time, but both DIMMs in a bank must be the same size.
- ▶ On the 1RX and 2RX models, if you want to install more than eight DIMMs, you must also replace the existing eight-DIMM memory board with the memory upgrade card, 73P9710.
- ▶ DIMMs sockets must be filled in a specific order to maintain balance across the two ports. This order is described in the *x365 Option Installation Guide*.
- ▶ If you want to install the full 32 GB, you will need to remove the existing 512 MB DIMMs and fully populate the memory board with 2 GB DIMMs.

## Memory mirroring

Memory mirroring is available on the x365 for increased fault tolerance. It requires the 16-DIMM memory board be installed.

The feature is roughly equivalent to RAID-1 in disk arrays, in that memory is divided in two ports and one port is mirrored to the other half (see Figure 8 on page 11). If 8 GB is installed, for example, then the operating system sees 4 GB once memory mirroring is enabled (it is disabled in the BIOS by default).

Memory mirroring is operating system independent, since all mirroring activities are handled by the hardware.

When memory mirroring is enabled, the data that is written to memory is stored in two locations. One copy is kept in the port 1 DIMMs, while a second copy is kept in the port 2 DIMMs. During the execution of the read command, the data is read from the DIMM with the least amount of reported memory errors through memory scrubbing, which is enabled with memory mirroring.

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

If memory scrubbing determines that the DIMM is damaged beyond use, read and write operations are redirected to the remaining good DIMM. Memory scrubbing then reports the damaged DIMM and the light path diagnostics display the error. If memory mirroring is enabled, then the mirrored copy of the data in the damaged DIMM is used until the system is powered down and the DIMM replaced. After the damaged DIMM is replaced, memory mirroring then copies the mirrored data back into the new DIMM.

Key configuration rules relating to memory mirroring:

- ▶ You must have the 16-DIMM board installed either as standard (3RX model) or as an option, part number 73P9710.
- ▶ Memory mirroring must be enabled in the BIOS (it is disabled by default).
- ▶ Both ports must have the same total amount of memory, and must have identical DIMMs. In other words, DIMMs must be installed in matched quads to support memory mirroring. Partial mirroring is not supported. See the *x365 Option Installation Guide* for information on the exact installation order required.

**Restriction:** Hot-add and hot-swap memory are not supported on the x365.

## PCI subsystem

As shown in Figure 3 on page 4, there are six full-length hot-swap PCI-X slots internal to the x365. These each support 3.3 V, 32-bit or 64-bit, PCI or PCI-X adapters. The slots support adapters of different speeds:

- ▶ Slot 1 is a 64-bit, 33 MHz slot. It contains the standard Remote Supervisor Adapter II.
- ▶ Slot 2 supports adapters up to 64-bit, 100 MHz. It also has an slot extension to support the ServeRAID 6i.
- ▶ Slot 3 supports adapters up to 64-bit, 133 MHz.
- ▶ Slot 4 supports adapters up to 64-bit, 133 MHz.
- ▶ Slot 5 supports adapters up to 64-bit, 133 MHz.
- ▶ Slot 6 supports adapters up to 64-bit, 133 MHz.

Slots 3–6 are on separate PCI-X buses so the adapter speeds they support are not affected by other devices. Slot 2 shares its PCI-X bus with the onboard Ethernet and SCSI controllers, so careful selection of the adapter in this slot is important, as it will affect the performance of these onboard devices.

See “x365 PCI slots” on page 14 for details on what adapters are supported.

The PCI subsystem also supplies these I/O devices:

- ▶ LSI Logic 1030 dual Ultra320 SCSI with integrated RAID-1
- ▶ Broadcom dual port 5704 10/100/1000 Ethernet
- ▶ Remote Supervisor Adapter II installed in Slot 1, which also supplies video
- ▶ EIDE interface
- ▶ Serial port
- ▶ Three USB ports (one on the front panel, two on the rear)

**Notes:** There is no parallel port on the x365. For parallel port connections, use the NetVista USB Parallel Printer Cable, part number 19K4164.

The Remote Supervisor Adapter II supplies the video for the server. The video subsystem is powered by an ATI Radeon chip with 8 MB of video memory.

With the addition of an RXE-100 Remote Expansion Enclosure, you can connect an additional six or 12 PCI-X adapters to the x365. See “RXE-100 Remote Expansion Enclosure” on page 18 for details.

## SCSI subsystem

The SCSI subsystem comprises a LSI Logic 1030 Dual Ultra320 SCSI controller. This allows for independent control of both channels. The Ultra320 ports are attached to an internal HDD backplane and an external connector at the rear of the system (see Figure 4 on page 5). The integrated SCSI controller with RAID capabilities supports only RAID level-1. The SCSI configuration is achieved using the LSI Logic Configuration Utility program.

ServeRAID adapters can also be installed and used to control the internal hot-swap drives or external disk enclosures such as the EXP300. The adapters currently supported are:

- ▶ ServeRAID-6i (must be installed in slot 2)
- ▶ ServeRAID-6M
- ▶ ServeRAID-4Lx
- ▶ ServeRAID-4Mx

The ServeRAID-4H is also supported but only for external disk enclosures, as the adapter with internal connectors is too high for the 3U server chassis.

See ServerProven for the latest list of supported adapters:

<http://www.pc.ibm.com/us/compat/controllers/matrix.html>

For a comparison of features of members of the ServeRAID family, see:

<http://www.redbooks.ibm.com/abstracts/tips0054.html>

**Note:** LSI RAID level-1 is not compatible with ServeRAID RAID level-1. Therefore, if the operating system is installed on drives connected to the integrated LSI chipset, the operating system will need to be reinstalled should the customer decide to move to a ServeRAID adapter. For more information, see:

<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-46174>

## x365 PCI slots

As shown in Figure 3, there are five full-length 64bit PCI-X slots and one full-length PCI slot internal to the x365.

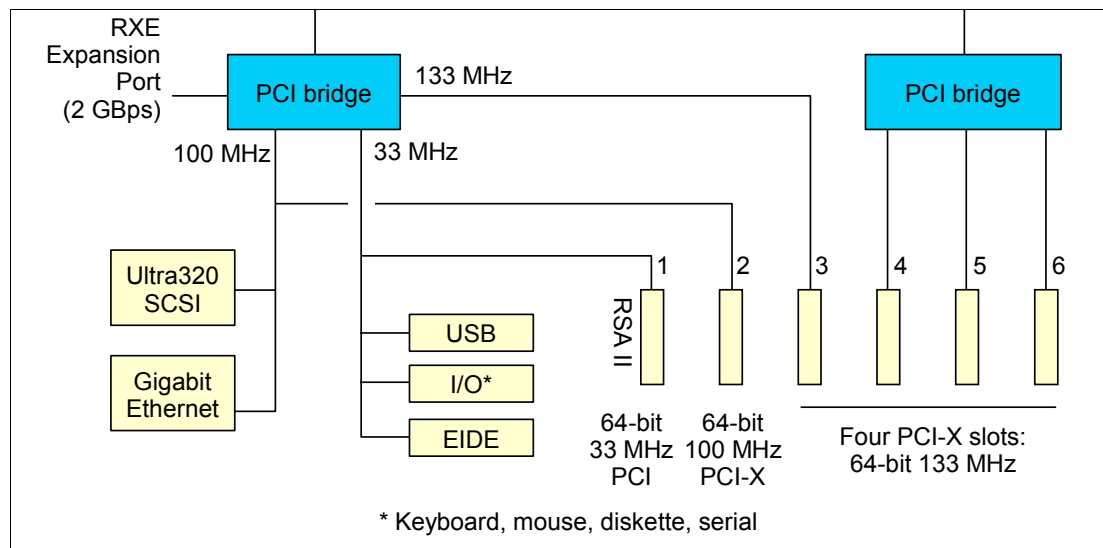


Figure 9 xSeries 365 system block diagram (partial)

These six slots are implemented as follows:

- ▶ Slot 1 is a 64-bit, 33 MHz PCI slot. It contains the standard Remote Supervisor Adapter II, which is a 32-bit 66 MHz adapter (and runs at 33 MHz in this slot). Slot 1 shares a PCI bus with devices such as USB, CD-ROM, diskette, mouse, and keyboard controllers.

- ▶ Slot 2 supports adapters up to 64-bit, 100 MHz. This slot supports the ServeRAID-6i adapter. (It is also the only slot that supports the ServeRAID-6i.) Slot 2 shares a PCI-X bus with the LSI SCSI controller and the Broadcom Ethernet controller.
- ▶ Slot 3 supports adapters up to 64-bit, 133 MHz. It does not share the bus with any devices. If you plan to install the IXA adapter (for connectivity to an IBM @server iSeries server), it must be installed in this slot.
- ▶ Slot 4 supports adapters up to 64-bit, 133 MHz. It does not share the bus with any devices.
- ▶ Slot 5 supports adapters up to 64-bit, 133 MHz. It does not share the bus with any devices.
- ▶ Slot 6 supports adapters up to 64-bit, 133 MHz. It does not share the bus with any devices.

Certain restrictions apply to adapters in these slots:

- ▶ Slot 2 will operate at 100 MHz or the rated speed of the installed adapter, whichever is slower. However, if you install an adapter that runs less than 100 MHz (for example, a 66 MHz card), it will adversely affect the performance of the SCSI and Ethernet controllers since these devices are on the same bus.
- ▶ Slots 3–6 can accept adapters rated at speeds ranging from 33 MHz to 133 MHz. The slot speed will slow down to the speed of the adapter. For example, if the adapter is a 66 MHz card, the bus will operate at 66 MHz.
- ▶ 32-bit adapters can be installed in any of the slots and will run in 32-bit mode. For slot 2, a 32-bit adapter and the 64-bit SCSI and Ethernet controllers can coexist on the same bus. The 32-bit adapter will run in 32-bit mode, and the 64-bit devices will run in 64-bit mode.
- ▶ The SCSI and Ethernet controllers on the bus with slot 2 are PCI-X devices. You can install a PCI (not PCI-X) adapter, but the entire bus, including the SCSI and Ethernet controllers, will operate in PCI mode with a loss in performance.

**Important:** If you install a PCI or slow PCI-X adapter in slot 2, you will impact the performance of the SCSI and Ethernet subsystems.

Note also that you cannot hot-add a PCI adapter into slot 2, because the bus is running in PCI-X mode and cannot switch to PCI mode without rebooting.

- ▶ The iSeries IXA adapter is a double thickness card. If you do plan to install the IXA adapter in slot 3, you will not be able to use slot 4.

**Tip:** Take the time to understand these rules and to select the best slots for your adapters. Incorrect choices can result in a loss of PCI adapter performance.

In summary, Table 4 lists the supported adapter speeds. Take into account the speed reductions when there are two adapters installed in a bus, as described above.

Table 4 Supported adapter speeds in each slot

Slot	Width (bits)	Supported adapter speed (MHz)
1	32 or 64	33 MHz (the standard 32-bit 66 MHz Remote Supervisor Adapter II operates in 32-bit 33 MHz mode)
2	32 or 64	33 or 66 or 100 (installing 33 and 66 adapters will affect performance of the SCSI and Ethernet controllers)
3	32 or 64	33, 66, 100 or 133
4	32 or 64	33, 66, 100 or 133

Slot	Width (bits)	Supported adapter speed (MHz)
5	32 or 64	33, 66, 100 or 133
6	32 or 64	33, 66, 100 or 133

Other configuration information:

- ▶ Video adapters are not supported.
- ▶ The PCI slots supports 3.3 V adapters only. 5 V adapters are not supported.
- ▶ The ServeRAID 4H adapter is not supported for internal drives because the adapter is too high to fit in the 4U server when a cable is attached to its internal connector.
- ▶ The system scans PCI-X slots to assign system resources. The system attempts to start the first device found. The search order is:
  - a. CD-ROM
  - b. Diskette drive
  - c. Integrated dual Gigabit Ethernet controller
  - d. Integrated SCSI devices
  - e. Internal PCI slot 1
  - f. Internal PCI-X slots (in the order 2, 3, 4, 5, 6)

If an RXE-100 is attached, the order is:

- a. CD-ROM
  - b. Diskette drive
  - c. Integrated dual Gigabit Ethernet controller
  - d. Integrated SCSI devices
  - e. Internal PCI slot 1
  - f. Internal PCI-X slots (in the order 2, 3, 4, 5, 6)
  - g. RXE-100 slots (A1, A2, A3, A4, A5, A6, B1, B 2, B3, B4, B5, B6)
- ▶ IBM Integrated xSeries Adapter (IXA) is supported for high-speed host attachment to iSeries.




Figure 10 The IBM Integrated xSeries Adapter

The following configuration points apply when installing this adapter:

- The adapter is only supported in slot 3 of the x365.
- The adapter takes up the space of two slots, so slot 4 is unusable.
- The IXA is a 66 MHz PCI adapter.



- Jumper 26 on the I/O planar must be moved from position 1-2 to 2-3. See the *xSeries 365 Option Installation Guide* for details.

See the following for more information about the adapter and integration with iSeries servers:

<http://www.ibm.com/servers/eserver/iseries/windowsintegration/>

## Broadcom dual Gigabit Ethernet controller

The x365 offers a dual Gigabit Ethernet controller integrated standard in the system. The x365 includes one dual-port Broadcom BCM5704 10/100/1000 BASE-T MAC (Media Access Controller) on the PCI-x 64-bit 100 MHz bus. The BCM5704 has the following features:

- ▶ Supports full and half-duplex performance at all speeds (10/100/1000 Mbps, auto negotiated).
- ▶ Includes integrated on-chip memory for buffering data transmissions to ensure the highest network performance.
- ▶ Includes dual onboard RISC processors for advanced packet parsing.
- ▶ Backward compatibility with today's 10/100 network.

The Broadcom controller also includes software support for failover, layer-3 load balancing, and comprehensive diagnostics.

Category 5 or better Ethernet cabling is required with RJ-45 connectors. If you plan to implement a Gigabit Ethernet connection, ensure that your network infrastructure is capable of the necessary throughput to match the server's I/O capacity.

## Remote Supervisor Adapter II

The x365 includes as standard the Remote Supervisor Adapter II service processor. This adapter provides the same functionality as the Remote Supervisor Adapter II option available for other xSeries servers. Key features of the Remote Supervisor Adapter II over the original Remote Supervisor Adapter include:

- ▶ System-independent graphical console redirection
  - Built-in video compression hardware eliminates drivers.
  - Graphics response, up to five times faster than with its predecessor, 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 or operating system installation and updates
    - Single server deployment
  - 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.
- Two serial ports supported; the extra serial port can be used to redirect the console of an additional device.
- 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 xSeries servers.
  - Secure Sockets Layer (SSL) encrypts the data transmitted between LDAP server and the Remote Supervisor Adapter II
- ▶ Investment protection
  - Supports IBM Interconnect (RS-485) connection with other xSeries servers.
  - Integrates with IBM Director and Director Agent.

The card has the following physical attributes:

- ▶ Half-length, 32-bit, up to 66 MHz PCI adapter (when installed in the x365, the adapter runs at 33 MHz)
- ▶ IBM ASIC with integrated PowerPC 405 core executing at 200 MHz
- ▶ 16 MB SDRAM and 4 MB flash ROM
- ▶ External cable with single high-density connector, dual RJ-48 ports, and two serial ports
- ▶ External connectors:
  - One high-density connector supporting two RS232 (9-pin) serial ports and the ASM interconnect
  - Analog video (system video)
  - 10/100 Base T Ethernet Port (RJ-45)
- ▶ Two status LEDs (externally visible on the bracket) and two status LEDs on the Ethernet port

## RXE-100 Remote Expansion Enclosure

The x365 can be connected to a single RXE-100 (8684-1RX) to provide an additional six or 12 PCI-X slots to the server. In addition, two x365 servers can share the slots of RXE-100—six slots each.

The RXE-100 has six 133 MHz 64-bit PCI-X slots as standard and can accept adapters with speeds ranging from 33 MHz to 133 MHz. With the optional six-slot expansion kit (part number 31P5998) installed, the RXE-100 has 12 slots. Each set of six adapter slots is divided into three buses of two slots each, as shown in Figure 11 on page 19.




Figure 11 RXE-100 PCI-X expansion board (6 slots)

For each of the three buses (A, B, C), one of the following can be installed:

- ▶ One 64-bit PCI-X 133 MHz adapter (in the odd-numbered slot), running at up to 133 MHz
- ▶ Two 64-bit PCI-X 133 MHz adapters running at up to 100 MHz
- ▶ Two 64-bit PCI or PCI-X, 33 or 66 MHz adapters

These slots can accept adapters rated at speeds ranging from 33 MHz to 133 MHz. When deciding which adapters to put in which slots, consider the following:

- ▶ Each adapter has a maximum rated speed and each bus also has a maximum rated speed.
- ▶ Installed adapters will operate at the slowest of three speeds:
  - The rated speed of adapter 1 in the bus
  - The rated speed of adapter 2 in the bus
  - The rated speed of the bus
- ▶ 32-bit adapters can be installed in any of the slots and will run in 32-bit mode. 32-bit and 64-bit adapters can coexist in 64-bit slots on the same bus. The 32-bit adapters will run in 32-bit mode, and the 64-bit adapters will run in 64-bit mode.
- ▶ When installing a 133 MHz PCI-X adapter, remember that it must be installed in the first or odd-numbered slot in the bus (that is, in slots 1, 3, or 5).
- ▶ Like the x365, a PCI-X and a PCI adapter can be installed in slots on the same bus in the RXE-100. However, these two adapters will both operate in PCI mode.

In addition, if you have a PCI-X adapter installed, you cannot hot-add a PCI adapter to the same bus. This is because with just the PCI-X adapter installed, the bus is running in PCI-X mode, and you cannot hot-add a PCI adapter into a bus that is in PCI-X mode.

- ▶ With Windows NT 4.0 Enterprise Edition, certain token-ring adapters do not work in some slots in the RXE-100. See RETAIN tip H175383 for more information:

<http://www.pc.ibm.com/qtechinfo/MIGR-42139.html>

## Cabling

The cabling used to connect the RXE-100 to the x365 is shown in Figure 12. The figure shows both an RXE-100 with six PCI-X slots installed and one with 12 slots installed. The

difference is the addition of a small loop between two of the management ports on the RXE-100.




Figure 12 Connecting the x365 to an RXE-100 (left: six slots in the RXE-100; right: 12 slots in the RXE-100)

There are two types of cables used to connect the RXE-100 to the x365:

- ▶ Remote I/O cable, for data

This cable connects from the x365 RXE Expansion Port to the RXE-100. Two lengths are available:

- 3.5 m Remote I/O cable kit (part number 31P6102)
- 8 m Remote I/O cable kit (part number 31P6103)

The RXE-100 ships with a 3.5 m Remote I/O cable to connect the unit to the x365. This cable is long enough when the devices are in the same rack as each other. For installation in an adjacent rack, use the optional 8 m Remote I/O cable kit.

- ▶ Interconnect management cable, for control

The RXE-100 also includes a 3.5 m interconnect management cable, which is used to connect the RXE Management Port on the x365 to the RXE Management A (In) Port on the RXE-100.

Two lengths are available:

- 3.5 m interconnect management cable kit (part number 31P6087)
- 8 m interconnect management cable kit (part number 31P6088)

**Important:** This cable has standard RJ-45 connectors on each end, but it is not the same wiring as a standard Ethernet network cable, nor a crossover cable. Ensure that the proper cables listed above are used for this connection.

If the RXE-100 has the second set of six PCI slots installed, use the short interconnect management cable (supplied with the PCI slot option kit) to connect Management A (out) Port to Management B (in) Port.

The 8 m interconnect management cable is suitable for inter-rack configurations.

**Tip:** Power to the RXE-100 is controlled by the x365, via the interconnect management cable and under the control of the Remote Supervisor Adapter.

### Sharing the RXE-100 with two x365 configurations

The x365 will support sharing of the RXE-100. The RXE-100 must have 12 PCI-X slots installed, and each x365 will have access to six slots. Connections are shown in Figure 13.

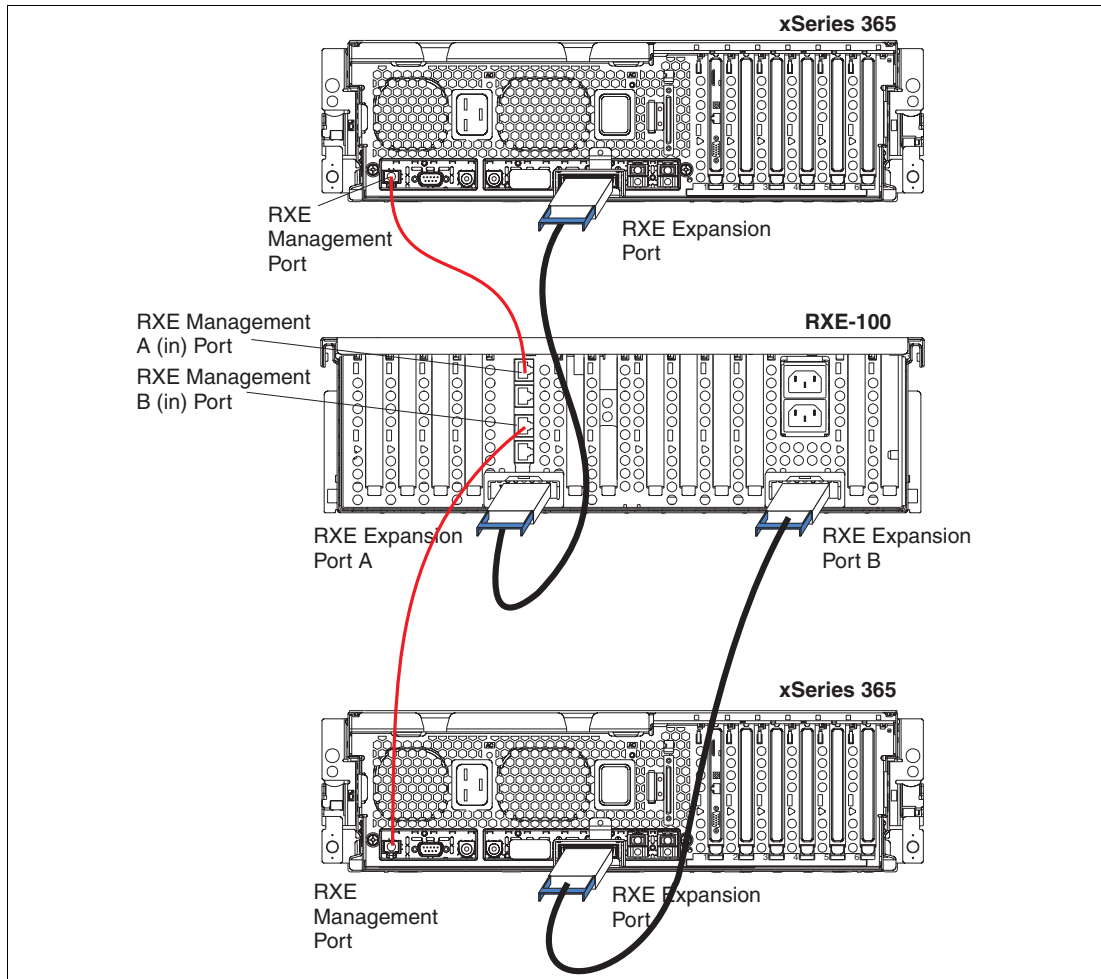


Figure 13 Sharing an RXE-100 between two x365 configurations

The data cable connects to the RXE-100 using the RXE Management B (in) Port. The short cable that is used in the non-shared 12-slot configuration is not used in this configuration.

This configuration requires additional cables be ordered:

- ▶ One remote I/O cable, either 3.5 m (part 31P6102) or 8 m (part 31P6103)
- ▶ One Interconnect management cable, either 3.5 m (31P6087) or 8 m (31P6088)

**Restriction:** There are no supported configurations for sharing Remote I/O enclosures with other servers, such as the x360, x440, x445, x450 or x455.

# Redundancy

The x365 has the following redundancy features to maintain high availability.

- ▶ Six hot-swap multi-speed fans.

Six, hot-swap, multispeed fans to provide cooling redundancy and enable individual fan replacement without powering down the server. In the event of a fan failure, the other fans 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.

- ▶ One or two 950 W hot-swap power supplies with separate power cords.

As per Table 1 on page 2, the 1RX models have one power supply standard and the 2RX and 3RX models have two power supplies standard. With either 110 V or 220 V power, the second power supply provided redundancy in case the first supply fails.

- ▶ Six hot-swap hard disk drive bays with onboard support for RAID-1.

Together with the integrated LSI Logic 1030 dual Ultra320 SCSI, they can be configured to create a RAID-1 disk array for the operating system. Various ServerRAID adapters are also supported to form other RAID levels.

- ▶ 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 memory mirroring, as described in “System memory” on page 8.

The layout of the front panel of the x365, showing the location of the drive bays, power supplies, and fans, is shown in Figure 14.




Figure 14 Redundancy features of the x365

## 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 x365. This panel slides out from the front of the server to view all light path diagnostics-monitored server subsystems. In the event that maintenance is then required, the customer can slide the server out from the rack and, using the LEDs, find the failed or failing component.

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




Figure 15 Light path diagnostic panel

Note that the pair of 7-segment LEDs on the top of the swing-out unit are development debug tools only.

## Operating system support

Table 5 on page 24 lists the supported operating systems for the x365 and the level of support for Hyper-Threading technology provided by the operating system. For the latest support information, see ServerProven:

<http://www.pc.ibm.com/us/compat/nos/matrix.shtml>

In the column titled Hyper-Threading Support:

- ▶ *None* indicates that the operating system does not recognize the logical processors that Hyper-Threading enables.
- ▶ *Yes* indicates that the operating system recognizes the logical processors and can execute threads on them but is not optimized for Hyper-Threading. From a licensing perspective, if Hyper-Threading is enabled, the operating system will need to be licensed for twice the number of physical processors to take full advantage of the processors' capabilities.
- ▶ *Optimized* indicates that the operating system recognizes the logical processors and the operating system code has been designed to take full advantage of the technology. From a

licensing perspective, the logical processors do not count towards the number of processors for which an operating system is licensed.

- *Ignores* indicates that the operating system recognizes the logical processors when Hyper-Threading is enabled, but the operating system ignores them.

Table 5 x365 operating system support

Operating system <sup>1</sup>	Release	Support <sup>2</sup>	Hyper-Threading
Windows 2000 Server	SP3	Yes	Yes
Windows 2000 Advanced Server	SP3	Yes	Yes
Windows Server 2003, Standard Edition	Initial	Yes	Optimized
Windows Server 2003, Enterprise Edition	Initial	Yes	Optimized
NetWare	6.5	Yes	Optimized
Red Hat Enterprise Linux AS	2.1	Yes	Yes
Red Hat Enterprise Linux	3.0	Planned	
SuSE Linux Enterprise Server	8.0	Yes	Yes
SuSE Linux Enterprise Server	9.0	Planned	
VMware ESX Server	2.0	Yes <sup>3</sup>	Ignores <sup>4</sup>

Notes for Table 5:

1. While operating systems may support eight-way or larger systems, scalability is a function of both the operating system and the application/workload. Applications must be specifically designed to take advantage of larger SMP systems.
2. Some operating systems are supported at general availability (GA)—these are marked with “Yes”. Those operating systems marked with “Planned” are planned to be supported after GA.
3. Virtual machines running on VMware ESX Server 2.0 will only be dual processor capable if the VMware ESX Virtual SMP module is also purchased.
4. Hyper-Threading can be enabled. However, VMware ESX Server 2.0 will *not* make use of any logical processors.

## The team that wrote this Redpaper

This Redpaper was produced by a technical specialist at the International Technical Support Organization, Raleigh Center.

**David Watts** is a Consulting IT Specialist at the International Technical Support Organization in Raleigh. He manages residencies and produces IBM Redbooks on hardware and software topics related to IBM xSeries systems and associated client platforms. He has authored over 35 IBM Redbooks, Redpapers and Solution Assurance Product Review Guides; his most recent books include *Implementing Systems Management Solutions using IBM Director*. He has a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for over 15 years. He is an IBM @server Certified Specialist for xSeries and an IBM Certified IT Specialist.



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

Dan Hurlimann, Senior Engineer, xSeries development, Austin  
Milan Miklos, WW Service Planning & Support, Raleigh  
Glenn Myrto, xSeries Mechanical Development, Raleigh  
Richard Rudd, x365 World Wide Product Manager, Raleigh  
Ori Schafman, xSeries Technical Specialist, Israel



# 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 illustrates 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. 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 IBM's application programming interfaces.

This document created or updated on January 29, 2004.



Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:  
[ibm.com/redbooks](http://ibm.com/redbooks)
- ▶ Send your comments in an Internet note to:  
[redbook@us.ibm.com](mailto:redbook@us.ibm.com)
- ▶ Mail your comments to:  
IBM Corporation, International Technical Support Organization  
Dept. HZ8 Building 662  
P.O. Box 12195  
Research Triangle Park, NC 27709-2195 U.S.A.

## Trademarks

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

@server™  
@server™  
ibm.com®  
iSeries™  
xSeries®  
Chipkill™

IBM®  
NetVista™  
Notes®  
PowerPC®  
Predictive Failure Analysis®  
Redbooks™

RETAIN®  
ServerProven®  
ServeRAID™  
Redbooks(logo) ™

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

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