



IBM Power 570 now supports the AIX, IBM i, and Linux operating systems

Description	2
Publications	17
Services	17
Technical information	17
IBM Electronic Services	21
Terms and conditions	22
Prices	22
Order now	33

At a glance

The Power 570 model 9117-MMA includes:

- Support for IBM i operating system and its I/O
- Machine type/model upgrades from models 9406-570 and 9406-MMA
- Hot-node Add / Cold-node Repair
- SAS RAID adapter support (#5902)

For ordering, contact:

Your IBM representative, an IBM Business Partner, or the Americas Call Centers at

800-IBM-CALL Reference: YE001

Overview

New features of the Power 570 model 9117-MMA include:

- Support for the IBM i operating system (formerly i5/OS™) in addition to the already supported AIX® and Linux™ operating systems.
- Machine type/model upgrades from the 9406-570 and the 9406-MMA, preserving the system serial number assigned as an IBM System i™ product.
- 12X Channel remote I/O drawers as features of the Power 570 server. Machine type/model 7314-G30 can be ordered as feature 5796. The 7314-G30 I/O drawer may be migrated, but new I/O drawers should be ordered as features.
- Alternate SAS controller for three of six internal Disk Slots (#5909) offers a single-slot split SAS backplane feature by combining the internal cable and the SAS controller function in the same slot in the CEC enclosure (P1-C3).
- PCI-X DDR Dual — X4 3 Gb SAS RAID Adapter (#5902), an excellent solution for high-performance applications. Two features are required and support redundant paths with proper cabling to the SAS drives. This dual adapter configuration provides mirrored write cache data and mirrored RAID parity footprints between the adapters for superior availability.
- DDR2 memory features 4496, 4497, 4499, and 7892 from the p5-570 can now be migrated during a model upgrade from the p5-570 to the Power 570 and are available for purchase with the feature 5621 processor.
- Power 570 systems now support adding an additional CEC enclosure (node) to a system, without powering down the system (Hot-node Add). The additional Power 570 enclosure would be ordered as a

system upgrade (MES order) and added to the original system while operations continued. The additional resources of the newly added enclosure can then be assigned to existing applications or new applications as required. This capability is available to existing 9117-MMA systems after upgrading to firmware level EM320_051, or later. To take advantage of this new feature, the system must have spare, unused connectors on the attached Service Interface Cable and space available in the rack immediately below the system. If a new Service Interface Cable is required, the system will need to be powered down to install the new cable.

- Power 570 systems that have experienced a failure and auto recovered by rebooting after Guarding Out the failed CEC enclosures can now have the deactivated enclosure removed so it can be repaired without powering down the system a second time. Once repaired, the enclosure can be reintegrated back into the active system without powering down the system (Cold-node Repair). The resources of the repaired CEC enclosure can then be reassigned to existing applications or new applications as required. This capability is available to all 9117-MMA systems after upgrading to firmware level EM320_051, or later.
- For Power 570 systems running IBM i partitions, a new Capacity Back-Up feature (#4991) is available that supports migrating the IBM i License entitlement from a primary system to a backup system on a temporary basis. For more information, visit

<http://www.ibm.com/systems/power/hardware/cbu>

Key prerequisites

Refer to the Software requirements section.

Planned availability date

May 6, 2008, except features 2146 and 2147, which have an availability date of April 8, 2008

Description

IBM Power 570

The POWER6™ processors in Power 570 model 9117-MMA are 64-bit processor cores. They are configured as dual cores on a single chip module, with 32 MB of L3 cache, 8 MB of L2 cache, and 12 DDR2 memory DIMM slots. The POWER6 processor is available at frequencies of 3.5 GHz, 4.2 GHz, and 4.7 GHz. The POWER6 DDR2 memory uses a new memory architecture to provide greater bandwidth and capacity. This enables operating at a higher data rate for large memory configurations. Each POWER6 processor card can support up to 12 DDR2 DIMMs running at speeds up to 667 MHz. A full system can contain up to 768 GB of memory in a 16-core system.

Available features within the system enclosures include:

- 4U 19-in rack-mount system enclosure
- One to four system enclosures; 16U maximum system size
- Two processor sockets per enclosure
- Six hot-swappable 3.5-in SAS disk bays per enclosure
- One hot-plug, slim-line media bay per enclosure (optional)
- Redundant hot-swap ac power supplies (N+1) in each enclosure
- Choice of integrated (HEA) I/O options — one per enclosure
 - 2-port 1 Gigabit Virtual Ethernet
 - 4-port 1 Gigabit Virtual Ethernet
 - 2-port 10 Gigabit Virtual Ethernet Short Run (SR)
- Two USB ports per enclosure
- Four HMC ports per system (two per enclosure) to support the attachment of two HMCs
- POWER6, 64-bit, 3.5 GHz, dual core processor (#5620):

- Twelve DDR2 POWER6 Memory DIMM sockets per processor card
- Two-, four-, eight-, twelve-, or sixteen-core configurations
- L2 cache: 4 MB per core, 8 MB per dual core
- L3 cache: 32 MB per dual core
- 2 GB to 384 GB of POWER6 DDR2 memory
- POWER6, 64-bit, 4.2 GHz, dual core processor (#5622):
 - Twelve DDR2 POWER6 Memory DIMM sockets per processor card
 - Two-, four-, eight-, twelve-, or sixteen-core configurations
 - L2 cache: 4 MB per core, 8 MB per dual core
 - L3 cache: 32 MB per dual core
 - 2 GB to 768 GB of POWER6 DDR2 memory
- POWER6, 64-bit, 4.2 GHz, dual core processor (#5621) (available for upgrade systems only):
 - Eight DDR2 DIMM sockets per processor card
 - Four-, eight-, twelve-, or sixteen-core configurations
 - L2 cache: 4 MB per core, 8 MB per dual core
 - L3 cache: 32 MB per dual core
 - 2 GB to 256 GB DDR2 memory (same memory features as 9117-570)
- POWER6, 64-bit, 4.7 GHz, dual core processor (#7380):
 - Twelve DDR2 POWER6 Memory DIMM sockets per processor card
 - Two-, four-, eight-, twelve-, or sixteen core configurations
 - L2 cache: 4 MB per core, 8 MB per dual core
 - L3 cache: 32 MB per dual core
 - 2 GB to 768 GB of POWER6 DDR2 memory
- POWER6 DDR2 Memory DIMMs
 - 0/4 GB (4 X 1 GB) POWER6 DDR2 Memory, 667 MHz
 - 0/8 GB (4 X 2 GB) POWER6 DDR2 Memory, 667 MHz
 - 0/16 GB (4 X 4 GB) POWER6 DDR2 Memory, 533 MHz
 - 0/32 GB (4 X 8 GB) POWER6 DDR2 Memory, 400 MHz
- Carry-over DDR2 Memory DIMMs (requires processor #5621):
 - 2 GB (4 X 0.5 GB) DDR2 Memory, 533 MHz (#7892)
 - 4 GB (4 X 1 GB) DDR2 Memory, 533 MHz (#7893)
 - 8 GB (4 X 2 GB) DDR2 Memory, 533 MHz (#7894)
 - 16 GB (4 X 4 GB) DDR2 Memory, 533 MHz (#4497)
 - 16 GB (4 X 4 GB) DDR2 Memory, 400 MHz (#4499)
 - 4/8 GB (4 X 2 GB) DDR2 Memory, 533 MHz (#4495)
 - 8/16 GB (4 X 4 GB) DDR2 Memory, 533 MHz (#4496)
- Seven I/O expansion slots per enclosure (28 max per system)

Slot ID	Adapter Type	Slot Size
P1-C1	PCIe 8X	Full length
P1-C2	PCIe 8X	Full length
P1-C3	PCIe 8X	Full length
P1-C4	PCI-X 2.0 DDR	Full length
P1-C5	PCI-X 2.0 DDR	Full length
P1-C6/P1-C8	PCIe 8X / GX+	Short form factor

Up to 16 PCIe 8X adapters
 Up to 8 GX+ Adapters
 Up to 8 PCI-X DDR Adapters

- Dynamic LPAR support
- PowerVM (optional):
 - Micro-Partitioning™ (up to 10 partitions per processor, 160 per system)
 - VIOS
 - Automated CPU and memory reconfiguration support for dedicated and shared processor logical partition (LPAR) groups
 - Support for manual provisioning of resources partition migration (with PowerVM — Enterprise Edition)
- Optional HACMP™ for AIX 5L™ support for nearly continuous operation
- Optional Cluster 1600 support with Cluster Systems Management software
- Remote I/O drawer and tower support:
 - Up to 20 I/O drawers on a RIO-2 interface (7311-D11)
 - Up to 48 I/O drawers or towers on a RIO-2 interface (7311-D20 or #0595, #5790, #0588, #5088, #5094, #5096, #5294, #5296)
 - Up to 32 I/O drawers on a 12X channel interface (7314-G30 or #5796)
 - Up to 110 SAS DASD I/O drawers on SAS PCI controllers (#5886)
 - Up to 60 DASD expansion drawers in i5/OS partitions (#5786)
 - Up to 40 DASD expansion drawers in AIX or Linux partitions (7031-D24)

Remote I/O drawer availability

PCI/SCSI Disk Expansion Drawer (#0595)

(i5/OS partition only)

The PCI/SCSI Disk Expansion Drawer (#0595) is a 5 EIA unit drawer and mounts in a 19-in rack. Feature 0595 is 24 inches long and can weigh up to 101 lb. The high-density expansion drawer provides additional adapter slots and SCSI disk slots as remote I/O. There are seven hot-swap PCI-X 64-bit, 133 MHz, 3.3 V I/O slots, and 12 optional hot-swap disk drive bays. Feature 0595 has redundant power and cooling. The fans, power supplies, and PCI adapters are top-accessible while the disk drives are front-accessible for easy service and maintenance. Feature 0595 attaches to a host system CEC enclosure with a RIO-2 adapter in a GX slot.

TotalStorage® EXP24 Disk Drawer (#5786)

(i5/OS partition only)

The TotalStorage EXP24 (#5786) is a 4 EIA unit drawer and mounts in a 19-in rack. The drawer is 27 inches long and can weigh up to 120 lb.

The front of the IBM TotalStorage EXP24 Ultra320 SCSI Expandable Storage Disk Enclosure has slots for up to 12 disk drives organized in two SCSI groups of up to six drives. The rear also has slots for up to 12 disk drives organized in two additional SCSI groups of up to six drives plus slots for the four SCSI interface cards. Each SCSI drive group can be connected by either a Single Bus Ultra320 SCSI Repeater Card (#5741) or a Dual Bus Ultra320 SCSI Repeater Card (#5742), allowing a maximum of eight SCSI connections per TotalStorage EXP24. Feature 5786 is delivered with three cooling fans and two power supplies to provide redundant power and cooling. Feature 5786 attaches to a host system CEC enclosure or to a remote I/O drawer with an Ultra320 SCSI adapter.

PCI Expansion Drawer (#5790)

(i5/OS partition only)

The PCI Expansion Drawer (#5790) is a 4 EIA unit tall drawer and mounts in a 19-in rack. Feature 5790 is 8.6 inches wide and takes up half the width of the 4 EIA rack space. Feature 5790 requires the use of a feature 7307 or 7311 drawer mounting enclosure. The 4 EIA tall

enclosure can hold up to two feature 5790 drawers mounted side by side in the enclosure. The drawer is 28 inches long and can weigh up to 37 lb. The PCI Expansion Drawer has six 64-bit, 3.3 V PCI-X slots that use blind-swap cassettes and support hot plugging of adapter cards. The drawer includes redundant hot-plug power and cooling. Feature 5790 attaches to a host system CEC enclosure with a RIO-2 adapter in a GX slot.

PCI DDR 12X Expansion Drawer (#5796)

The PCI-DDR 12X Expansion Drawer (#5796) is a 4 EIA unit tall drawer and mounts in a 19-in rack. Feature 5796 is 8.8 inches wide and takes up half the width of the 4 EIA rack space. Feature 5796 requires the use of a feature 7314 drawer mounting enclosure. The 4 EIA tall enclosure can hold up to two feature 5796 drawers mounted side by side in the enclosure. The drawer is 31.5 inches deep and can weigh up to 44 lb. The PCI-DDR 12X Expansion Drawer has six 64-bit, 3.3 V PCI-X DDR slots running at 266 MHz that use blind-swap cassettes and support hot plugging of adapter cards. The drawer includes redundant hot-plug power and cooling. You must select one of the two available interface adapters for use in the feature 5796 drawer. The Dual-Port 12X Channel Attach Adapter — Long Run (#6457) or the Dual-Port 12X Channel Attach Adapter Short Run (#6446). The adapter selection is based on how close the host system or the next I/O drawer in the loop is physically located. Feature 5796 attaches to a host system CEC enclosure with a 12X adapter in a GX slot.

EXP 12S SAS Drawer (#5886)

The EXP 12S SAS drawer (#5886) is a 2 EIA unit tall drawer and mounts in a 19-in rack. The drawer is 20.12 inches long and can weigh up to 40 lb without SAS disks. The EXP 12S SAS drawer has twelve 3.5-in SAS disk slots with redundant data paths to each slot. The drawer supports redundant hot-plug power and cooling and redundant hot-swap SAS expanders with Enclosure Services Manager (ESM). Each ESM has an independent SCSI Enclosure Services (SES) diagnostic processor. Feature 5886 attaches to a host system CEC enclosure or a remote I/O drawer with a SAS adapter in a PCI-X or PCIe slot.

7311-D20 Rack Mounted High Density Expansion Drawer

(AIX/Linux Partitions only)

The 7311-D20 Expansion Drawer is a 4 EIA unit drawer and mounts in a 19-in rack. It is 24 inches long and can weigh up to 101 lb. The high-density expansion drawer provides additional adapter slots and SCSI disk slots as remote I/O. There are seven hot-swap PCI-X 64-bit, 133 MHz, 3.3 V I/O slots and 12 optional hot-swap disk drive bays. The drawer includes redundant power and cooling. The fans, power supplies, and PCI adapters, are top-accessible while the disk drives are front-accessible for easy service and maintenance. The D20 attaches to a host system CEC enclosure with a RIO-2 adapter.

19-in racks

The Model 9117-MMA and its remote I/O drawer features are designed to mount in the 7014-T00 or the 7014-T42 rack. These are both built to the 19-in EIA standard. When you order a new 9117 system, you can also order the appropriate 7014 rack model with the system hardware on the same initial order. IBM is making the racks available as features of the 9117-MMA when you order additional I/O drawer hardware for an existing system (MES order). The rack feature number should be used if you want IBM to integrate the newly ordered I/O drawer in a 19-in rack before shipping the MES order.

1.8-meter rack (#0551)

The 1.8-meter rack (#0551) is a 36 EIA unit tall rack. The rack that is delivered as feature 0551 is the same rack delivered when you order the 7014-T00 rack; the included features may be different. Some features that are delivered as part of the 7014-T00 must be ordered separately with the feature 0551. Order the feature 0551 only when required to support rack integration of MES orders prior to shipment from IBM.

2.0-meter rack (#0553)

The 2.0-meter rack (#0553) is a 42 EIA unit tall rack. The rack that is delivered as feature 0553 is the same rack delivered when you order the 7014-T42 rack; the included features may be different. Some features that are delivered as part of the 7014-T42 must be ordered separately with the feature 0553. Order the feature 0553 only when required to support rack integration of MES orders prior to shipment from IBM.

Capacity on Demand

Several types of Capacity on Demand (CoD) are optionally available on the Power 570 server to help meet changing resource requirements in an on demand environment by using resources

installed on the system but not activated.

Capacity Upgrade On Demand

Capacity Upgrade on Demand (CUoD) allows you to purchase additional permanent processor or memory capacity and dynamically activate the capacity when needed.

Trial CoD

Trial CoD is a no-additional-charge feature that allows you up to 30 days use of available inactive processor and memory resource. It is intended to be used as a benchmarking aid to determine how your applications will perform with the addition of the available inactive resource in your system. Trial CoD is enabled by registering at the CoD Web site and electronically receiving an activation key. To request either a standard or an exception trial, visit

https://www-912.ibm.com/tcod_reg.nsf/TrialCod?OpenForm

A standard trial provides up to 30 contiguous days use of up to two additional processors or up to 4 GB of memory, depending on the installed inactive hardware on the system. One standard trial is included with each new system. A standard processor or memory trial is renewed for the system when a permanent activation of the respective type is purchased on an MES order.

An exception trial offers up to 30 contiguous days use of all inactive processors or memory in the system. One processor and memory exception trial is supplied with each new system. Exception trial use is one time only and is not renewed when additional resources are purchased.

On/Off CoD

On/Off CoD enables processors or memory to be temporarily activated in full-day increments as needed. Charges are based on usage reporting collected monthly. Processors and memory can be activated and turned off an unlimited number of times, whenever you want additional processing resources. This offering gives a system administrator an interface at the HMC to manage the activation and deactivation of resources. A monitor that resides on the server logs the usage activity. You must send this usage data to IBM monthly. A bill is then generated based on the total amount of processor and memory resources utilized, in increments of processor and memory (1 GB) days. Before using temporary capacity on your server, you must enable your server. To do this, order an enablement feature (MES only) and sign the required contracts.

If a Power 570 server uses the IBM i operating system in addition to any other supported operating system on the same server, the client must inform the sales team placing the billing feature order which operating system caused the temporary On/Off CoD processor use so that the correct feature can be used for billing.

For more information regarding registration, enablement, and usage of On/Off CoD, visit

<http://www.ibm.com/systems/power/hardware/cod>

Utility CoD

Utility CoD autonomically delivers additional processor performance on a temporary basis within the shared processor pool. Utility CoD enables you to place a quantity of inactive processors into the server's default shared processor pool, which then become available to the pool's resource manager. When the server recognizes that the combined processor utilization within the shared pool exceeds 100% of the level of base (purchased/active) processors assigned across uncapped partitions, then a utility CoD processor minute is charged and this level of performance is available for the next minute of use.

If additional workload requires a higher level of performance, the system will automatically allow the additional Utility CoD processors to be used and the system automatically and continuously monitors and charges for the performance needed above the base (permanent) level. Registration and usage reporting for utility CoD is made using a public Web site and payment is based on reported usage. Utility CoD requires PowerVM (#7942 or #7995) to be active and Utility CoD is only available on model MMA of machine type 9117.

If a Power 570 server uses the IBM i operating system in addition to any other supported operating system on the same server, the client must inform the sales team placing the billing feature order which operating system caused the temporary Utility CoD processor use so that the correct feature can be used for billing.

For more information regarding registration, enablement, and use of Utility CoD, visit

<http://www.ibm.com/systems/power/hardware/cod>

Software licensing

For software licensing considerations with the various CoD Offerings, refer to the latest revision of the Capacity on Demand Planning Guide. Visit

<http://www.ibm.com/systems/power/hardware/cod>

Services

IBM Server Product Services offers implementation and migration services to help you put your IBM Power Systems server quickly into your production environment in order to support your business applications. These services include in-depth planning sessions to help ensure the end result is in line with your requirements. A variety of product services are available to suit your needs. The goal is to continually enhance these offers to provide you with a comprehensive selection of services. To find out what IBM can do for you, visit

<http://www.ibm.com/services/servers>

Model upgrades

You can upgrade the 9406-570, 9406-MMA, or 9117-MMA systems with IBM POWER5™, POWER5+™, or POWER6 processors to the Power 570 with POWER6 processors. For upgrades from POWER5 or POWER5+ processor-based systems, IBM will install new CEC enclosures to replace the enclosures that you currently have. You must return your current CEC enclosures to IBM in exchange for the financial considerations identified under the applicable feature conversions for each upgrade. Several of the parts in your current system can be moved to the new system after it is installed.

Clients taking advantage of the model upgrade offer from a POWER5 or POWER5+ processor-based system are required to return to IBM all CEC enclosures, bezels, and internal planars from the CEC enclosures that are part of the serial number system being upgraded. Any feature for which a feature conversion is used to obtain a new part must be returned to IBM also. Clients may keep and reuse any PCI adapters, SCSI disks, memory, or media devices that were not involved in a feature conversion transaction.

Some memory DIMMs used in the POWER5+ (p5-570) memory subsystem may have compatibility issues when used in a POWER6 (Power 570) memory subsystem. IBM will replace these DIMMs when a system is upgraded from POWER5+ processors to POWER6 processor feature number 5621. If IBM determines some of the DIMMs need to be replaced, IBM will exchange them at no additional charge for systems being upgraded to POWER6 processor feature 5621. When you are ready to place an order to upgrade your POWER5+ processor-based system to a POWER6 processor-based system with processor feature 5621, contact your local IBM Service office and request a service call for ECA840 on your POWER5+ system. IBM Service will assist you in obtaining a snap file from your system and in submitting the snap file to IBM for analysis. This action is required to allow your upgrade order to POWER6 processor feature 5621 to be scheduled for shipment from IBM. If you are upgrading your POWER5+ system to processor feature 5622, your memory does not migrate and ECA840 is not required.

You can upgrade the following:

- 9117-570 systems with POWER5+ processor feature numbers 7830, 7832, 7833, or 7834 to a 9117-MMA with processor feature number 5622.
- 9117-570 systems with POWER5+ processor feature numbers 7782 or 8338 to a 9117-MMA with processor feature numbers 5621 or 5622. You can order processor feature 5621 only during the initial upgrade MES order. Systems built initially with processor feature 5621 can be expanded with processor feature 5622. The minimum supported configuration when upgrading to processor feature 5621 is a 4-core system.
- 9406-570 systems with processor feature number 8961, 8971, or 8338 to a 9117-MMA with processor feature number 7380.
- 9406-MMA systems with processor feature number 7380 to a 9117-MMA with processor feature number 7380.

The following features, if present on the current system, can be moved to the new system.

- Operator panel.
- PCI adapters with cables: A maximum of two PCI-X adapters in each of the four system

enclosures. Remaining PCI-X adapters must be moved to attached remote I/O drawers or towers or replaced with PCIe adapters available for the 9117-MMA. Recommended I/O drawers are features 0595, 5790, and 5796.

- Line cords, keyboards, and displays.
- Advanced POWER™ Virtualization (#7942) (now called PowerVM — Standard Edition).
- SCSI DASD cannot be used in the MMA system enclosures but can be used in remote I/O drawers or towers attached to the model MMA. Recommended I/O drawers are feature 5786 or 7031-D24, or 7311-D20.
- DDR2 memory, if upgrading to the feature number 5621 processor.
- Power supplies, if upgrading to the feature number 5621 processor.
- Most I/O drawers and towers, if present as a feature on the 9406 system.
- Racks, if present as a feature on the 9406 system.

Upgrade considerations

The following should be considered when planning an upgrade to a 9117-MMA.

Each model 570 has six PCI-X adapter slots in each CEC enclosure. The model MMA has two PCI-X and four PCIe adapter slots for each CEC enclosure. When upgrading, PCIe adapters can be purchased to provide the function of the extra PCI-X adapters, or you can migrate the PCI-X adapters to a remote I/O drawer attached to the model MMA.

The model MMA supports the 7413-G30 PCI-X (12X channel) I/O drawer or the 7311-D11 or 7311-D20 (RIO-2) I/O drawers. In addition, the model MMA supports migration of the following I/O drawers when you upgrade from the 9406 machine type.

- PCI-X Expansion Unit in Rack (#0588)
- PCI-X Tower Unit in Rack (#0595)
- PCI-X Expansion Unit (#5088)
- PCI-X Expansion Tower (#5094)
- PCI-X Expansion Tower (no disk) (#5096)
- 1.8M I/O Tower (#5294)
- 1.8M I/O Tower (no disk) (#5296)
- TotalStorage EXP24 Disk Drawer (#5786)
- PCI Expansion Drawer (#5790)
- PCI-DDR 12X Expansion Drawer (#5796)

The model 570 has two native RIO ports on each CEC enclosure that support one RIO-2 loop per enclosure. The model MMA does not have native RIO ports but adds a second GX+ slot to each CEC enclosure and can support two GX+ adapters per enclosure.

The model MMA supports only the new SAS DASD hard disks internally. The older SCSI DASD hard disks can be attached to the model MMA but must be located in a remote I/O drawer. Feature 5786 or 7031-D24 or 7311-D20 are compatible with the DASD carriers used on the model 570 DASD files. The DASD files from the model 570 can be moved directly into these I/O drawers.

If there is a machine type 7310 HMC attached to the model 570, it can be moved to the model MMA, but the HMC firmware must be updated to Licensed Machine Code V7, or later, to accommodate the model MMA improvements. The firmware upgrade must be completed before attaching the 7310 HMC to the 9117-MMA.

For processor feature 5621, different memory size or frequency features cannot be mixed on the same processor card. All of the memory features on a single processor card must be the same size in gigabytes when fully active and have the same frequency. Features 7894 and 4495 can be mixed on the same processor card because feature 4495 is 8 GB when fully active. Features 4497 and 4496 can be mixed on the same processor card because feature 4496 is 16 GB when fully active. The two processor cards in a single CEC enclosure can have different memory features installed. Processor cards in different enclosures in the same system can have different

memory features installed.

For POWER5 or POWER5+ processor-based systems that have the On/Off CoD function enabled, you must reorder the On/Off enablement features (#7951 and #7954) when placing the upgrade MES order for the new Power 570 system to keep the On/Off CoD function active. To initiate the model upgrade, you need to remove the On/Off enablement features from the configuration file before the MES order is started. Any temporary use of processors or memory owed to IBM on the existing system must be paid before installing the new Power 570 model MMA.

Reserve CoD features are not supported on the model MMA.

Feature 8017 supports migration of memory activation code 7663 during the initial order and build of the upgrade MES MMA order. When you order an upgrade with processor feature 5621 and migrate CoD memory activations from the p5-570 to the Power 570, you should order feature 8017 in a quantity equal to the quantity of feature 7663 that is currently active on your p5-570 system.

Feature 8018 supports migration of the Advanced Power Virtualization feature 7942 (now called PowerVM — Standard Edition) during the initial order and build of the upgrade MES MMA order. You can add feature 8018 to your upgrade orders in a quantity not to exceed the quantity of feature 7942 previously obtained for the system being upgraded. Migrate the 7942 feature to the new configuration report in a quantity that equals 8018. Additional 7942 features can be ordered during the upgrade.

Machine type 9406 contained processor Micro-Partitioning as part of the base product definition. To maintain this function on the 9117 machine type, order PowerVM feature 7942 in a quantity at least equal to the number of active processors in the 9117 system. If feature 7942 was ordered on the 9406 machine type, the quantity purchased can be migrated to the 9117 machine type at no additional charge.

When upgrading from the 9406 machine type, many of the features that will migrate during this upgrade will be renamed and given a new feature number. The function in the system will remain the same and there is no cost associated with the renaming of the feature. The marketing configurator will identify and rename these features for you. It is important that the configuration file imported to the marketing configurator be accurate with regard to the content of the system being upgraded.

When upgrading from the 9406 machine type, it is important to place the order in the IBM order system only when you are ready to proceed with the upgrade. During the upgrade order process, the complete order will be split into two parts: a record purpose only (RPO) part, to rename many of the features in the system being upgraded, and a hardware MES upgrade order part, to deliver the new hardware. There is no additional charge for the RPO part of this upgrade order, and it will proceed quickly upon submission of the order. Once the upgrade order is placed, there can be no additional MES orders placed against this serial number until the upgrade system is installed. If the upgrade order is cancelled, a follow-up RPO order must be placed to reverse the renaming that was done as part of the original RPO order. No additional MES orders can be placed against the serial number system until this follow-up RPO order is complete.

Customers upgrading to a 9117-MMA from a 9406-570 or 9406-MMA system that is still under warranty will continue to have the warranty service level of the 9406 for the remainder of the 9406 one-year warranty period. Because the 9117 has 9x5 NBD base warranty coverage instead of the 9406 24x7 SD base warranty coverage, you will be required to sign a separate warranty service upgrade contract to continue 24x7 SD base warranty coverage. Contact your Business Partner or IBM representative for more details.

PowerVM

PowerVM Editions are available as a hardware feature (#7942 for Standard, #7995 for Enterprise). Select the feature that provides the level of virtualization required for your application.

- Partition Mobility, available only with PowerVM — Enterprise (#7995), enables you to move a running AIX or Linux LPAR from one physical server to another with no downtime if both servers are using POWER6 processors. This is a significant advancement in Power Systems virtualization technology and can assist you with performing the following common business needs while keeping the business up and running:
 - Evacuating a system before performing scheduled maintenance
 - Moving workloads across a pool of different physical resources as business needs shift
 - Moving workloads off underutilized machines so that they can be powered off to save on

energy and cooling costs

- Micro-Partitioning support for a single core being shared by up to 10 logical partitions.
- Virtual I/O Server is a single-function appliance that resides in a POWER5 or POWER6 processor-based partition. It facilitates the sharing of physical I/O resources between client partitions (AIX V5.3 or later, IBM i V6.1, or Linux) within the server. VIOS provides virtual SCSI targets and shared Ethernet adapter (SEA) virtual I/O to client LPARs.
- Virtual SCSI (VSCSI) enables the sharing of physical storage adapters (SCSI and Fibre Channel) and storage devices (disk and optical) between LPARs.
- Virtual networking: SEA enables connectivity between internal and external virtual LANs (VLANs); virtual Ethernet enables high-speed connections between partitions.
- Lx86 supports running most x86 Linux applications within Linux partitions (refer to the PowerVM Lx86 — Limitations section.)

PowerVM Editions:

- PowerVM — Standard (#7942) supports up to 10 partitions per core, HMC management, VIOS, Lx86, and multiple shared processor pools.
- PowerVM — Enterprise (#7995) adds support for Live Partition Mobility for AIX and Linux partitions.

Other PowerVM technologies include:

- Workload Partitions (WPARs) provide isolated instances on top of a single AIX 6.1 image.
- Live Application Mobility available with WPAR Manager provides the movement of a running AIX application from one server to another.
- System Planning Tool simplifies the process of planning and deploying Power Systems LPARs and virtual I/O.
- HMC provides management of virtualized systems.

PowerVM Lx86 — Limitations: PowerVM Lx86 runs most x86 Linux applications, but PowerVM Lx86 cannot run applications that:

- Directly access hardware devices (for example, 3D graphics adapters)
- Require nonstandard kernel module access or use kernel modules not provided by the Linux operating system distribution
- Do not use only the Intel® IA-32 instruction set architecture as defined by the 1997 Intel Architecture Software Developer's Manual consisting of Basic Architecture (order number 243190), Instruction Set Reference Manual (order number 243191), and the System Programming Guide (order number 243192) dated 1997
- Do not run correctly on Red Hat Enterprise Linux 4 starting with version 4.3 or Novell SUSE Linux Enterprise server (SLES) 9 starting with version SP3 or Novell SLES 10;
- Use real-mode or 16-bit instructions
- Are x86 Linux specific system administration or configuration tools

A more complete list of qualifications can be found at

<http://www.ibm.com/systems/power/software/virtualization/editions/lx86/qual.html>

Capacity Backup (CBU) Offering (applies to i5/OS only)

The Power 570 systems' CBU designation can help meet your requirements for a second system to use for backup, high availability, and disaster recovery. It enables you to temporarily transfer i5/OS processor license entitlements and 5250 Enterprise Enablement entitlements purchased for a primary machine to a secondary CBU-designated system. Temporarily transferring these resources instead of purchasing them for your secondary system may result in significant savings. Processor activations cannot be transferred.

The CBU specify feature 4891 is available only as part of a new server purchase or during an MES upgrade from an existing CBU system to a 9117-MMA. Certain system prerequisites must be met and system registration and approval are required before the CBU specify feature can be applied on a new server.

Standard i5/OS terms and conditions do not allow either i5/OS processor license entitlements or 5250 OLTP (Enterprise Enablement) entitlements to be transferred permanently or temporarily. These entitlements remain with the machine they were ordered for. When you register the association between your primary and on-order CBU system, you must agree to certain terms and conditions regarding the temporary transfer.

After a CBU system designation is approved and the system is installed, you can temporarily move your optional i5/OS processor license entitlement and 5250 Enterprise Enablement entitlements from the primary system to the CBU system when the primary system is down or while the primary system processors are inactive. The CBU system can then better support failover and role swapping for a full range of test, disaster recovery, and high availability scenarios. Temporary entitlement transfer means that the entitlement is a property transferred from the primary system to the CBU system and may remain in use on the CBU system as long as the registered primary and CBU system are in deployment for the high availability or disaster recovery operation.

The primary system for a Power 570 server can be:

- 9119-FHA
- 9406-595
- 9117-MMA
- 9406-MMA
- 9406-570

These systems have i5/OS software licenses with an i5/OS P30 software tier or higher. The primary machine must be in the same enterprise as the CBU system.

Before you can temporarily transfer i5/OS processor license entitlements from the registered primary system, you must have more than one i5/OS processor license on the primary machine and at least one i5/OS processor license on the CBU server. An activated processor must be available on the CBU server to use the transferred entitlement. You may then transfer any i5/OS processor entitlements above the minimum one, assuming the total i5/OS workload on the primary system does not require the i5/OS entitlement you would like to transfer during the time of the transfer. During this temporary transfer, the CBU system's internal records of its total number of i5/OS processor license entitlements are not updated, and you may see i5/OS license noncompliance warning messages from the CBU system. Such messages that arise in this situation do not mean you are not in compliance.

Before you can temporarily transfer 5250 entitlements, you must have more than one 5250 Enterprise Enablement entitlement on the primary server and at least one 5250 Enterprise Enablement entitlement on the CBU system. You may then transfer the entitlements that are not required on the primary server during the time of transfer and that are above the minimum of one entitlement.

For example, if you have an eight-core Power 570 as your primary system with four i5/OS processor license entitlements (three above the minimum) and two 5250 Enterprise Enablement entitlements (one above the minimum), you can temporarily transfer up to three i5/OS entitlements and one 5250 Enterprise Enablement entitlement. During the temporary transfer, the CBU system's internal records of its total number of i5/OS processor entitlements is not updated, and you may see i5/OS license noncompliance warning messages from the CBU system.

If your primary or CBU machine is sold or discontinued from use, any temporary entitlement transfers must be returned to the machine on which they were originally acquired.

For CBU registration and further information, visit

<http://www.ibm.com/systems/power/hardware/cbu>

Reliability, availability, and serviceability (RAS)

The reliability of the IBM Power 570 starts with components, devices, and subsystems that are designed to be fault-tolerant. During the design and development process, subsystems go through rigorous verification and integration testing processes. During system manufacturing,

systems go through a thorough testing process to help ensure high product quality levels.

The Power 570 L3 cache and system memory offers error checking and correcting (ECC) fault-tolerant features. ECC is designed to correct environmentally induced, single-bit, intermittent memory failures and single-bit hard failures. With ECC, the likelihood of memory failures is substantially reduced. ECC also provides double-bit memory error detection that helps protect data integrity in the event of a double-bit memory failure.

AIX supports disk mirroring and disk controller duplexing while Linux supports DASD mirroring (RAID 1) through software. Additionally, some hardware RAID adapters are supported on Linux.

When using AIX, the Journaled File System, also known as JFS or JFS2, helps maintain file system consistency and reduces the likelihood of data loss when the system is abnormally halted due to a power failure. JFS, for 32-bit kernels, and JFS2, for 64-bit kernels, are designed to substantially reduce or eliminate fragmentation.

Hot-node Add / Cold-node Repair

- Power 570 systems support adding an additional CEC enclosure (node) to a system, without powering down the system (Hot-node Add). The additional Power 570 enclosure would be ordered as a system upgrade (MES order) and added to the original system while operations continued. The additional resources of the newly added enclosure can then be assigned to existing applications or new applications as required. This capability is available to existing 9117-MMA systems after upgrading to firmware level EM320_051, or later. To take advantage of this new feature, the system must have spare, unused connectors on the attached Service Interface Cable and space available in the rack immediately below the current system. If a new Service Interface Cable is required, the system will need to be powered down to install the new cable.
- Power 570 systems that have experienced a failure and auto-recovered by rebooting after Guarding Out the failed CEC enclosures can have the deactivated enclosure removed so it can be repaired without powering down the system a second time. Once repaired, the enclosure can be reintegrated back into the active system without powering down the system (Cold-node Repair). The resources of the repaired CEC enclosure can then be reassigned to existing applications or new applications as required. This capability is available to existing 9117-MMA systems after upgrading to firmware level EM320_051, or later.

Memory error-correction extensions

POWER6 memory has single-error-correct and double-error-detect ECC circuitry designed to correct single-bit memory failures. Double-bit detection is designed to help maintain data integrity by detecting and reporting multiple errors beyond what the ECC circuitry can correct. The memory chips are organized such that the failure of any specific memory module is designed to only affect a single bit within an ECC word (bit scattering), thus allowing for error correction and continued operation in the presence of a complete DRAM chip failure (Chipkill™ recovery).

POWER6 memory also uses memory scrubbing and dynamic bit steering, enabling correctable error thresholding to determine when available spare memory modules on each DIMM should be used to replace ones that have exceeded their threshold value.

Redundancy for array self-healing

Although the most likely failure event in a processor is a soft single-bit error in one of its caches, other events can occur, and they need to be distinguished from one another. For the L1, L2, and L3 caches and their directories, hardware and firmware keep track of whether permanent errors are being corrected beyond a set threshold. If exceeded, a deferred repair error log is created. Additional runtime availability actions, such as CPU vary off or L3 cache line delete, are also initiated.

L1 and L2 caches and L2 and L3 directories on the POWER6 chip are manufactured with spare bits in their arrays that can be accessed via programmable steering logic to replace faulty bits in the respective arrays. This is similar to the redundant bit steering employed in main storage as a mechanism designed to help avoid physical repair; it is also implemented in systems using the POWER6 processor. The steering logic is activated during processor initialization and is initiated by the built-in system-test (BIST) at power-on time.

When correctable error thresholds in the L3 cache exceed a set threshold, systems using the POWER6 processor invoke a dynamic L3 cache line delete function, which enables them to stop using bad cache and eliminates exposure to greater problems.

Fault monitoring functions

When a system using the POWER6 processor is powered on, BIST and power-on self-test

(POST) check processors, L3 cache, memory, and associated hardware required for proper booting of the operating system. If a noncritical error is detected or if errors occur in resources that can be removed from the system configuration, booting the system is designed to proceed to completion. The errors are subsequently logged in the system nonvolatile RAM.

Disk drive fault tracking can alert the system administrator of an impending disk failure before it affects customer operation.

On POWER6 processor-based servers running the AIX or Linux operating systems, hardware and software failures are recorded in the system log. An Error Log Analysis (ELA) routine analyzes the error, forwards the event to the Service Focal Point (SFP) application running on the HMC, and notifies the system administrator that it has isolated a likely cause of the system problem. The service processor event log also records unrecoverable checkstop conditions and forwards them to the SFP application and notifies the system administrator. After the information is logged, if the system is properly configured, a call home service request is initiated and the pertinent failure data with service parts information and part locations is sent to an IBM Service organization. Customer contact information and specific system-related data, such as the machine type, model, and serial number along with engineering data related to the failure, are sent to IBM service.

The call home feature enables IBM service representatives to preemptively bring the most-probable replacement parts when a service call is placed, reducing repair time.

Mutual surveillance

The service processor monitors the operation of firmware during the boot process and also monitors the Hypervisor for termination. The Hypervisor monitors the service processor and performs a reset/reload if it detects the loss of the service processor. If the reset/reload does not correct the problem, the Hypervisor notifies the operating system and the operating system can take appropriate action, including calling for service.

Environmental monitoring functions

POWER6 processor-based servers include a range of environmental monitoring functions.

- Temperature monitoring increases the fan speed rotation when ambient temperature is above the normal operating range or when a redundant fan fails.
- Temperature monitoring warns the system administrator of potential environmental-related problems (for example, air conditioning and air circulation around the system) so that appropriate corrective actions can be taken before a critical failure threshold is reached. It also performs an orderly system shutdown when the operating temperature exceeds the critical level.
- Fan speed monitoring provides a warning and increases fan speeds on the redundant fans to compensate for a detected fan failure. It can also initiate an orderly system shutdown if there is an additional cooling failure when the fans are in high-speed mode.
- Voltage monitoring includes a warning and an orderly system shutdown when the voltages are out of operational specification.

Availability enhancement functions

The POWER6 line of systems continues to offer and introduce significant enhancements designed to increase system availability.

POWER6 processor availability enhancements

One of the significant mainframe-inspired availability enhancements in systems with the POWER6 processor is the ability to perform processor instruction retry and alternate processor recovery. This significantly reduces exposure to both hard (logic) and soft (transient) errors in the processor core.

Soft failures in the processor core are transient (intermittent) errors, often due to cosmic rays or other sources of radiation, and generally are not repeatable. When an error is encountered in the core, the POWER6 processor will first automatically retry the instruction. If the source of the error is truly transient, the instruction will succeed and the system will continue as before. On previous IBM systems, this error would cause a checkstop.

With hard failures, true logical errors that are replicated each time the instruction is repeated, retrying the instruction does not help. The instruction will continue to fail. Systems with POWER6 processors can extract the failing instruction from the faulty core and retry it elsewhere in the system, after which the failing core is dynamically deconfigured and called out for replacement. The entire process is transparent to the partition owning the failing instruction. Systems with

POWER6 processors are designed to avoid what would have been a full system outage on earlier models.

POWER6 single processor checkstopping

Another major advancement in POWER6 processors is single processor checkstopping. The Power 570 can contain most processor checkstops to the partition that was using the processor at the time. This significantly reduces the probability of any one processor affecting total system availability.

POWER6 cache availability

The POWER line of servers continues to be at the forefront of cache availability enhancements. While L3 cache line delete (also called Pellston healing) was introduced with POWER4™ processors, the POWER6 processor line pioneers L2 cache line delete. If an uncorrectable error occurs in L2 or L3 cache, the system can dynamically remove the offending line of cache without requiring a reboot. In addition, the POWER6 processor-based server uses an inclusive L1/L2 cache design and a write-through cache policy on all levels. This helps ensure that data is written to main memory as soon as possible. POWER6 processors also continue to offer hardware-assisted memory scrubbing.

Special uncorrectable error handling

Special uncorrectable error handling, an IBM innovation introduced for POWER5 processors, prevents an uncorrectable error in memory or cache from immediately causing the system to terminate. Rather, the system tags the data and determines whether it will ever be used again. If the error is irrelevant, it will not force a checkstop.

PCI extended error handling

In systems using POWER6 processors, I/O drawer hardware, system firmware, and AIX interaction are designed to allow transparent recovery of intermittent PCI bus parity errors and graceful transition to the I/O device available state in the case of a permanent parity error in the PCI bus. This mechanism is called PCI extended error handling (EEH).

EEH-enabled adapters respond to a special data packet generated from the affected PCI slot hardware by calling system firmware, which examines the affected bus, enables the device driver to reset it, and continues without a system reboot. Currently, there is limited support for Linux, depending upon driver availability.

Predictive Failure Analysis® and dynamic component deal location

Servers with POWER processors have long been able to perform Predictive Failure Analysis on certain critical components such as processors and memory. When these components exhibit symptoms that would indicate a failure is imminent, the system can dynamically deallocate and call home about the failing part before the error becomes systematic. In many cases, this is transparent, especially if the system contains CUoD components. If no CUoD resources are available, the system will first attempt to reallocate resources in such a way that will avoid unplanned outages. If insufficient resources exist to maintain full system availability, these servers will attempt to maintain partition availability by user-defined priority.

Uncorrectable error recovery

The auto-restart (reboot) option, when enabled, can reboot the system automatically following an unrecoverable software error, software hang, hardware failure, or environmentally induced (ac power) failure.

Serviceability

The IBM Power 570 is designed with both IBM and customer serviceability in mind.

Advancements such as Guiding Light LED architecture are used to control a system of integrated LEDs that lead the individual servicing the machine to the correct part as quickly as possible.

With the Power 570, you can replace service parts (customer replaceable unit). To do this, the Power 570 uses Guiding Light LEDs to indicate the parts that need to be replaced.

An HMC attached to the Power 570 enables support personnel (with your authorization) to remotely log in to review error logs and perform remote maintenance if required.

The I/O device and adapter diagnostics consist of stand-alone diagnostics, which are loaded from the DVD-ROM drive, and online diagnostics. Online diagnostics, when installed, are resident with the AIX operating system on the disk or system. They can be booted in single-user mode (service mode), run in maintenance mode, or run concurrently (concurrent mode) with other applications. They have access to the AIX error log and the AIX configuration data.

- Service mode enables checking of system devices and features.
- Concurrent mode allows the normal system functions to continue while selected resources are being checked.
- Maintenance mode enables checking of devices and adapters.

Note: Because the 9117-MMA system has an optional DVD-ROM (#5756) and DVD-RAM (#5757), alternative methods for maintaining and servicing the system need to be available if the DVD-ROM or DVD-RAM is not ordered; an external Internet connection must be available to maintain or update system microcode to the latest required level.

Concurrent maintenance guided service procedures will continue to be supported by the Repair and Verify (R & V) component of the Service Focal Point application running on the HMC. Repair procedures that are not covered by the guided R & V component will be documented and available for display on any Web browser-enabled system as well as on the HMC. Instead of using the InfoCenter like the 9117-570 and the 9406-570, the Power 570 service procedures will be released through the System Support Site. Subscribe to the System Support Site to obtain the latest service procedure updates as they are released from IBM.

Error handling and reporting

In the unlikely event of system hardware or environmentally induced failure, system run-time error capture capability systematically analyzes the hardware error signature to determine the cause of failure. CEC and power and cooling subsystem events are analyzed concurrently with system operation by the processor runtime diagnostics (PRD) or by Hypervisor code. Results of these analyses are logged on the service processor and forwarded to both the operating system partitions and the SFP application on the HMC.

The SFP application contains logs of all detected and analyzed errors, detected either by the device drivers on the operating system for I/O-based events or by PRD or Hypervisor for runtime-detected errors in the remainder of the system.

With the call home function on the HMC, the system can automatically send out an alert via phone line to a pager or initiate a call for service when a service action point is reached.

A hardware fault also turns on the two attention indicators (one located on the front of the system unit and the other on the rear) to alert the user of an internal hardware problem. The indicator can also be turned on by the operator as a tool for system identification. The indicators will flash to signify identification but will remain on when an error condition occurs.

Service processor

The Power 570 service processor enables you to diagnose, check the status of, and sense the operational conditions of a system. It runs on its own power boundary and does not require resources from a system processor to be operational to perform its tasks.

The service processor supports surveillance of the connection to the HMC and to the system firmware (Hypervisor). It also includes several remote power control options, environmental monitoring (but only critical errors are supported under Linux), reset and boot features, and remote maintenance and diagnostic functions, including console mirroring. The service processors menus (ASMI) can be accessed concurrently with system operation allowing nondisruptive abilities to change system default parameters.

Call home

The call home application is available at no additional charge and is installed on every HMC. When it is properly configured, it helps IBM deliver the best maintenance service for your system.

Concurrent maintenance

The Power 570 continues to support concurrent repair of power, cooling, PCI adapters, media devices, and the operator panel. In addition, it continues to support concurrent firmware updates when possible. The determination of whether a firmware release can be updated concurrently is identified in the readme file released with the firmware.

IBM Electronic Services

IBM Electronic Services are designed to increase system availability and automatically manage system support.

IBM Electronic Service Agent™ and IBM Electronic Services Web Portal comprise the IBM Electronic Services solution. IBM is dedicated to providing fast, exceptional support to clients

who use IBM servers.

Electronic Service Agent automatically reports hardware events, such as system errors and performance issues. Systems enabled with this tool "call home" to IBM Support, 24 hours a day, 7 days a week — at no additional cost.

As part of an increased focus on better service, Electronic Service Agent is configured at installation time for every IBM Power 570. IBM systems support representatives configure the tool when the system is installed, at no charge. In support of this program, a new security white paper describes the secure data exchange between the HMC and the IBM Service Delivery Center (SDC) and the methods and protocols for this exchange. To read the white paper and prepare for Electronic Service Agent installation, visit the Reference Guide section at

<http://www.ibm.com/support/electronic>

Integrated with the HMC, and part of the AIX and i5/OS operating systems, Electronic Service Agent is designed to automatically and electronically report system failures and utilization issues to IBM, resulting in faster problem resolution and increased availability for your systems. System configuration and inventory information collected by Electronic Service Agent can also be viewed by IBM Support for faster problem resolution.

The Electronic Services Web Portal is a single Internet entry point that replaces the multiple entry points traditionally used to access IBM Internet services and support. This Web portal enables you to gain easier access to IBM resources for assistance in resolving technical problems. The newly improved My Systems and My Search functions make it even easier for you to track system inventory and find pertinent fixes.

Benefits

- **Higher availability:** Electronic Service Agent is designed to enhance your Warranty or Maintenance Agreement by offering faster hardware error reporting and uploading system information to IBM Support. This can translate to less wasted time monitoring the symptoms, diagnosing the error, and manually calling IBM Support to open a problem record. Around-the-clock monitoring and reporting means no more dependence on human intervention or off-hours customer personnel when errors are encountered in the middle of the night.
- **Security:** Electronic Service Agent is secure — monitoring, reporting, and storing the data at IBM. Electronic Service Agent securely transmits data either via the Internet or modem, and can be configured to communicate securely through gateways to give you a single point of exit from your site. Communication is one way — IBM cannot call into your system. System inventory information is stored in a secure database that is protected behind IBM firewalls. Your business applications or business data is never transmitted to IBM.
- **More accurate reporting:** Better reporting and automatic problem routing mean higher system availability and fix confidence. Because system information and error logs are automatically uploaded to the IBM Support center in conjunction with the service request, you do not need to find and send system information, decreasing the risk of misreported or misdiagnosed errors. Once inside IBM, problem error data is run through a data knowledge management system and knowledge articles are appended to the problem record.
- **Customized support:** Using the IBM ID you entered during activation, you can view system and support information in the My Systems and My Search sections of the Electronic Services Web Portal at

<http://www.ibm.com/support/electronic>

My Systems offers formatted reports of hardware and software installed using information collected from the systems by Electronic Service Agent. Reports are available for any system associated with your IBM ID. My Search combines the function of search and system information collected by Electronic Service Agent to deliver advanced searching of the IBM technical support knowledgebase, with results customized to your unique system.

For more information on how to take advantage the power of Electronic Services and for information about Electronic Service Agent for all your platforms, contact your IBM Systems Services representative or visit

<http://www.ibm.com/support/electronic>

Business Partner information

If you are a Direct Reseller - System Reseller acquiring products from IBM, you may link directly to Business Partner information for this announcement. A PartnerWorld ID and password are required (use IBM ID).

BP Attachment for Announcement Letter 108-259

<https://www.ibm.com/partnerworld/mem/sla.jsp?num=108-259>

Trademarks

i5/OS, System i, POWER6, Micro-Partitioning, AIX 5L, HACMP, POWER5, POWER5+, POWER, Chipkill, POWER4, and Electronic Service Agent are trademarks of International Business Machines Corporation in the United States or other countries or both.

AIX, TotalStorage, Predictive Failure Analysis, and xSeries are registered trademarks of International Business Machines Corporation in the United States or other countries or both.

Intel is a registered trademark of Intel Corporation.

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

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

Publications

No publications are shipped with these features.

Services

Global Technology Services

IBM services include business consulting, outsourcing, hosting services, applications, and other technology management.

These services help you learn about, plan, install, manage, or optimize your IT infrastructure to be an On Demand Business. They can help you integrate your high-speed networks, storage systems, application servers, wireless protocols, and an array of platforms, middleware, and communications software for IBM and many non-IBM offerings. IBM is your one-stop shop for IT support needs.

For details on available services, contact your IBM representative or visit

<http://www.ibm.com/services/>

For details on available IBM Business Continuity and Recovery Services, contact your IBM representative or visit

<http://www.ibm.com/services/continuity>

For details on education offerings related to specific products, visit

<http://www.ibm.com/services/learning/index.html>

Select your country, and then select the product as the category.

Technical information

Software requirements: If installing the AIX® operating system (one of these):

- AIX 5.2 with Technology Level 5200-10, or later
- AIX 5.3 with Technology Level 5300-06, or later

- AIX 6.1, or later

If installing the IBM i operating system (one of these):

- IBM i 5.4, with V5R4M5 machine code, or later
- IBM i 6.1, or later

If installing the Linux™ operating system (one of these):

- SUSE Linux Enterprise Server 10 SP1 for POWER™, or later
- Red Hat Enterprise Linux 4.5 for POWER, or later
- Red Hat Enterprise Linux 5.1 for POWER, or later

Notes:

- PowerVM (#7942 and #7995) are not supported on AIX 5L™ for POWER V5.2.
- Not all features available with the Power 570 are supported with each available operating system. Check the specific feature detail to identify which of the available operating systems are supported.

Limitations: Racks

The 36 EIA unit (1.8 meter) rack (#0551) and the 42 EIA unit (2.0 meter) rack (#0553) are available for order on MES upgrade orders only. For initial system orders, the racks should be ordered as machine type 7014, models T00 or T42.

The Power 570 consists of one to four CEC enclosures. Each enclosure occupies 4U of vertical rack space. The Power 570 is designed to be installed in a 7014-T00 or -T42 rack. An existing -T00 or -T42 rack can be used for the Power 570 if sufficient space and power are available. The Power 570 is not supported in the 7014-S11.

For Power 570 configurations with two, three, or four drawers, all drawers must be installed together in the same rack, in a contiguous space of 8U, 12U, or 16U within the rack. The uppermost enclosure in the system is the base enclosure. This enclosure will contain the active service processor and the operator panel, if an operator panel is present in the system. If a second CEC enclosure is part of the system, the backup service processor is contained in the second CEC enclosure. The service processor is a component of the Service Interface Card in these enclosures.

When a Power 570 system is installed in an 7014-T00 or -T42 rack or in the feature 0551 or 0553 rack that is ordered without a rack front door, a Thin Profile Front Trim Kit must be ordered for the rack. Order feature number 6246 for the 7014-T00 or feature 0551 rack and feature number 6247 for the 7014-T42 or feature 0553 rack.

The design of the Power 570 is optimized for use in a 7014-T00 or -T42 rack. Both the front cover and the processor flex cables occupy space on the front left side of an IBM 7014 rack that may not be available in typical non-IBM racks.

Acoustic Door features are available with the 7014-T00 and 7014-T42 racks to meet the lower acoustic levels identified in the specification section. The Acoustic Door feature can be ordered on new T00 and T42 racks or ordered for the T00 and T42 racks that you already own.

I/O drawers

To further reduce possible single points of failure, POWER6 based-technology implements enhanced disk storage configuration rules. IBM configuration tools and IBM technical support personnel do not support integrated cached disk controller configurations unless they have a protected write cache. Disk controllers with write cache must protect the cache by either pairing the disk controllers (write cache replication or IOA-level mirroring) or by using an auxiliary write cache IOA. This is true for all partitions in the Power 570 using any supported operating systems.

The remote I/O interfaces, HSL and HSL-2, on machine type 9406, are the same interfaces as RIO and RIO-2 respectively on the IBM Power 570. There is no difference between the HSL and RIO interfaces or the HSL-2 and RIO-2 interfaces.

The following list shows the remote I/O drawers that are supported or available on the 9117 machine type and the correct interface to use for each of the drawers.

Feature description	Order status	Interface
0588 PCI-X Exp Unit in Rack	Supported	RIO-2 (formerly HSL-2)
0595 PCI/SCSI Disk Exp Drwr	Available	RIO-2 (formerly HSL-2)
5088 PCI-X Exp Unit	Supported	RIO-2 (formerly HSL-2)
5094 PCI-X Exp Twr	Supported	RIO-2 (formerly HSL-2)
5096 PCI-X Exp Twr (no disk)	Supported	RIO-2 (formerly HSL-2)
5294 1.8M I/O Twr	Supported	RIO-2 (formerly HSL-2)
5296 1.8M I/O Twr (no disk)	Supported	RIO-2 (formerly HSL-2)
5786 EXP24 Disk Drwr	Available	SCSI Ultra320
5790 PCI Exp Drwr	Available	RIO-2 (formerly HSL-2)
5796 PCI-DDR 12X Exp Drwr	Available	12X
5886 Exp 12S SAS Drwr	Available	SAS
7031-D24/T24	Supported	SCSI Ultra320
7311-D11	Supported	RIO-2
7311-D20	Available	RIO-2
7314-G30	Supported	12X

7311-D20 and feature 0595 I/O drawers with RIO Ports, I/O Planar Riser Card (#6413) must be upgraded to RIO-2 Ports, I/O Planar Riser Card (#6417) before they can be attached to a Power 570 server.

Some I/O adapters supported in the 7311-D11 and the 7311-D20 I/O drawers when attached to a IBM POWER5. or POWER5+ processor-based Power Systems server will not be supported when these same I/O drawers are attached to a Power 570 server with POWER6 processors. For a complete list of supported adapters on the 7311, refer to the online Sales Manual for the I/O drawer.

The maximum number of attached remote I/O drawers depends on the number of CEC enclosures in the system and the I/O attachment type, as follows:

For the 7311-D11 RIO-2 attached I/O drawers:

- Systems with one CEC enclosure:
 - With two processors, the enclosure supports up to four I/O drawers.
 - With four processors, the enclosure supports up to eight I/O drawers.
- Systems with two CEC enclosures support up to 12 I/O drawers.
- Systems with three CEC enclosures support up to 16 I/O drawers.
- Systems with four CEC enclosures support up to 20 I/O drawers.

For the 7311-D20 and all feature number RIO-2 attached I/O drawers:

- Systems with one CEC enclosure:
 - With two processors, the enclosure supports up to six I/O drawers.
 - With four processors; the enclosure supports up to 12 I/O drawers.
- Systems with two CEC enclosures support up to 24 I/O drawers.
- Systems with three CEC enclosures support up to 36 I/O drawers.
- Systems with four CEC enclosures support up to 48 I/O drawers.

For 12X Host channel attached I/O drawers:

- Systems with one CEC enclosure:
 - With two processors, the enclosure supports up to four I/O drawers.
 - With four processors, the enclosure supports up to eight I/O drawers.
- Systems with two CEC enclosures support up to 16 I/O drawers.
- Systems with three CEC enclosures support up to 24 I/O drawers.
- Systems with four CEC enclosures support up to 32 I/O drawers.

It is recommended that any attached remote I/O drawers be located in the same rack as the

Power 570 server for ease of service, but they can be installed in separate racks if the application or other rack content requires it.

I/O drawers are connected to the adapters in the CEC enclosure with the following cables:

- Data transfer cables:
 - RIO-2 attach cables for RIO-2 I/O drawers
 - 12X cables for 12X Channel I/O drawers
- Power control cables

RIO-2 I/O drawers and 12X channel I/O drawers may not be mixed in the same remote I/O loop.

Remote I/O drawer cable connections are always made in loops to help protect against a single point-of-failure resulting from an open, missing, or disconnected cable. A Power 570 system with nonlooped configurations could experience degraded performance and serviceability. If a nonloop connection is detected, a problem is reported.

The first I/O drawer attached in any remote I/O drawer loop requires two data transfer cables. Each additional drawer in the loop (up to the maximum allowed) requires one additional data transfer cable.

The first I/O drawer attached to a host system requires two Power Control Cables. Each additional I/O drawer added to a host system requires one additional power control cable. Each host system has one power control loop. All I/O drawers attached to a system are included in the same power control loop. Power control cable loops are different in this regard from data transfer cable loops.

Dual-Port 12X Channel interface adapter options: Dual-Port 12X Channel Attach Adapter (#6446): Use the short run adapter for expansion I/O drawers located in close proximity to the host system or to other drawers in the I/O expansion loop. This adapter does not include signal repeaters.

Dual-Port 12X Channel Attach Adapter (#6457): Use the long run adapter for expansion I/O drawers located farther from the host system or other I/O drawers in the I/O expansion loop. This adapter includes signal repeaters to accommodate the longer cable lengths.

12X cable choice: Each 5796 drawer requires one Dual-Port 12X Channel Adapter, either Short Run (#6446) or Long Run (#6457). The choice of adapters is dependent on the distance to the next 12X Channel connection in the loop, either to another I/O drawer or the host system. The following table identifies the supported cable lengths for each 12X Channel Adapter. I/O drawers containing the Short Run adapter can be mixed in a single loop with I/O drawers containing the Long Run adapter. In this table, a "Yes" indicates that the 12X cable identified in that column can be used to connect the drawer configuration identified to the left. A "No" means it cannot be used.

	12X Cable Options			
	0.6 M (#1829) (1)	1.5 M (#1830)	3.0 M (#1840) (2)	8.0 M (#1834) (3)
5796 to 5796 w/12X Short Run adapter (#6446) in both drawers	Yes	Yes	No	No
5796 w/ 12X Short Run adapter (#6446) to 5796 w/ 12X Long Run adapter (#6457)	Yes	Yes	Yes	No
5796 to 5796 w/12X Long Run adapter (#6457) in both drawers	Yes	Yes	Yes	Yes
5796 w/12X Short Run adapter (#6446) to host system	No	Yes	Yes	No
5796 w/12X Long Run adapter (#6457)	No	Yes	Yes	Yes

to system

- (1) The 0.6M 12X cable (#1829) has very limited use due to its short length. It cannot be used to connect to a system drawer because of the short length. It is intended for use between two 5796 drawers mounted side by side in the same enclosure (#7314). It can also be used to connect between two modules located one beneath the other in a 7014 rack.
- (2) It is possible in some limited configurations to use the 3.0 M, 12X cable (#1840) to locate 5796 modules in adjacent racks. The cable length requires careful management of each drawer location within the rack. The best choice for connecting a G30 I/O Drawer in an adjacent rack is the 8.0 M, 12X cable (#1834).
- (3) The 8M 12X cable (#1834) is intended for use when connecting between two modules that are located in adjacent racks. This cable may not be connected to the 12X Short Run adapter (#6446).

Integrated I/O: Although each CEC enclosure is equipped with integrated system port (serial) external connectors, these ports do not function with the attachment of the required HMC.

Each CEC enclosure must contain one Virtual Ethernet (HEA) Integrated I/O port card (#5636, #5637, or #5639). This selection is available only when a CEC enclosure is first ordered from the factory.

Each system has two HMC ports on the Service Interface Card in each CEC enclosure. The HMC, however, must be attached to the Service Interface Card in the base CEC enclosure or the base and second CEC enclosure if a second CEC enclosure is part of the Power 570 system.

PowerVM Lx86:

PowerVM Lx86 runs most x86 Linux applications, but PowerVM Lx86 cannot run applications that:

- Directly access hardware devices (for example, 3D graphics adapters)
- Require nonstandard kernel module access or use kernel modules not provided by the Linux operating system distribution
- Do not use only the Intel® IA-32 instruction set architecture as defined by the 1997 Intel Architecture Software Developer's Manual consisting of Basic Architecture (order number 243190), Instruction Set Reference Manual (order number 243191), and the System Programming Guide (order number 243192) dated 1997
- Do not run correctly on Red Hat Enterprise Linux 4 starting with version 4.3 or Novell SUSE Linux Enterprise server (SLES) 9 starting with version SP3 or Novell SLES 10
- Use real-mode or 16-bit instructions
- Are x86 Linux specific system administration or configuration tools

A more complete list of qualifications can be found at

<http://www.ibm.com/systems/power/software/virtualization/editions/lx86/qual.html>

Security, auditability, and control

This product uses the security and auditability features of host software and application software.

The customer is responsible for evaluation, selection, and implementation of security features, administrative procedures, and appropriate controls in application systems and communications facilities.

IBM Electronic Services

IBM has transformed its delivery of hardware and software support services to help you achieve higher system availability. Electronic Services is a Web-enabled solution that offers an exclusive,

no-additional-charge enhancement to the service and support available for IBM servers. These services are designed to provide the opportunity for greater system availability with faster problem resolution and preemptive monitoring. Electronic Services comprises two separate, but complementary, elements: Electronic Services news page and Electronic Services Agent.

The Electronic Services news page is a single Internet entry point that replaces the multiple entry points traditionally used to access IBM Internet services and support. The news page enables you to gain easier access to IBM resources for assistance in resolving technical problems.

The Electronic Service Agent™ is no-additional-charge software that resides on your server. It monitors events and transmits system inventory information to IBM on a periodic, client-defined timetable. The Electronic Service Agent automatically reports hardware problems to IBM. Early knowledge about potential problems enables IBM to deliver proactive service that may result in higher system availability and performance. In addition, information collected through the Service Agent is made available to IBM service support representatives when they help answer your questions or diagnose problems. Installation and use of IBM Electronic Service Agent for problem reporting enables IBM to provide better support and service for your IBM server.

To learn how Electronic Services can work for you, visit

<http://www.ibm.com/support/electronic>

Terms and conditions

MES discount applicable: No. Equal to the volume commitment discount.

Field-installable features: Yes

Warranty period: This feature assumes the same warranty or maintenance terms as the machine in which it is installed for the full warranty or maintenance period announced for such machine.

Customers upgrading to a 9117-MMA from a 9406-570 or 9406-MMA system that is still under warranty will continue to have the warranty service level of the 9406 for the remainder of the 9406 one-year warranty period. Because the 9117 has 9x5 NBD base warranty coverage instead of the 9406 24x7 SD base warranty coverage, you will be required to sign a separate warranty service upgrade contract to continue 24x7 SD base warranty coverage. Contact your Business Partner or IBM representative for more details.

Customer setup: Yes, except for feature 5108

Licensed Machine Code: Same license terms and conditions as designated machine.

Prices

Type/Model conversions

From Type	Model	To Type	Model	Parts returned	Purchase price
9406	570	9117	MMA	Yes	\$9,195
9406	MMA	9117	MMA	Yes	0

The following are newly announced features on the specific models of the 9117 machine type:

Description	Machine Type	9117-MMA	Feature number	Purchase price	Initial / MES / Both / Support	RP CSU	MES
Specify Code for External High Speed Modem			0032	\$ 687	Both	Yes	No
Mirrored System Disk Level			0040	0	Both	Yes	No

Specify Code

Device Parity Protection- All Specify Code	0041	0	Both	Yes No
Mirrored System IOP Level Specify Code	0042	0	Both	Yes No
Mirrored System Bus Level Specify Code	0043	0	Both	Yes No
Device Parity RAID-6 All Specify Code	0047	0	Both	Yes No
RISC-to-RISC Data Migration	0205	0	Initial	Yes No
1Gbps Ethernet Specify	0226	0	Support	Yes No
CSC Specify	0275	0	Both	Yes No
Ext Tape Attached via #5736	0290	0	Both	Yes No
Specify Custom Data Protection	0296	0	Both	Yes No
Specify EXP24 Attach via Existing Controller	0302	0	MES	Yes No
Mirrored Level System Specify Code	0308	0	Both	Yes No
IPCS Extension Cables for NT	0325	144	Support	Yes No
RAID Hot Spare Specify	0347	0	Both	Yes No
V. 24/EIA232 6.1m (20-Ft) PCI Cable	0348	164	Both	Yes No
V. 24/EIA232 15.2m (50-Ft) PCI Cable	0349	229	Support	Yes No
V. 35 6.1m (20-Ft) PCI Cable	0353	164	Both	Yes No
V. 35 15.2m (50-Ft) PCI Cable	0354	229	Support	Yes No
V. 36 6.1m (20-Ft) PCI Cable	0356	164	Support	Yes No
X. 21 6.1m (20-Ft) PCI Cable	0359	491	Both	Yes No
X. 21 15.2m (50-Ft) PCI Cable	0360	229	Support	Yes No
X. 21 15.2m (50-Ft) PCI Cable	0365	262	Support	Yes No
V. 24/EIA232 24.4m (80-Ft) PCI Cable	0367	164	Both	Yes No
UPS Factory Integration Specify	0373	0	MES	N/A No
HMC Factory Integration Specify	0374	0	MES	N/A No
Display Factory Integration Specify	0375	0	MES	N/A No
Reserve Rack Space for UPS	0376	0	MES	N/A No
Reserve Rack Space for HMC	0377	0	MES	N/A No
Reserve Rack Space for Display	0378	0	MES	N/A No
MTM Upgrade Indicator	0395	0	MES	Yes No
512MB DDR Server Memory	0446	576	Support	Yes No

1GB DDR Server Memory	0447	1,532		Support	Yes No
IBM i 5.4 w/ V5R4M5 Specify Code	0533	0		Both	Yes No
IBM i 6.1 Specify Code	0534	0		Both	Yes No
19 inch, 1.8 meter high rack	0551	3,500		MES	Yes No
19 inch, 2.0 meter high rack	0553	4,745		MES	Yes No
19 inch, 1.3 meter high rack	0555	2,617		Support	Yes No
PCI-X Expansion Unit in Rack	0588	15,709	\$200	Support	Yes No
PCI/SCSI Disk Expansion Drawer	0595	9,419	229	Both	Yes No
Rack Filler Panel Kit	0599	98		Both	Yes No
#5094 Equivalent	0694	0		Support	Yes No
#5096 Equivalent	0696	0		Support	Yes No
Load Source Not in CEC	0719	0		Both	Yes No
Load Source in #0595	0720	0		Both	Yes No
Specify Load Source in #5786	0725	0		Both	Yes No
Specify #5886 Load Source placement	0727	0		Both	N/A No
#4319 Load Source Specify	0830	0		Support	Yes No
#4326 Load Source Specify	0834	0		Support	Yes No
#4327 Load Source Specify	0835	0		Both	Yes No
#4328 Load Source Specify	0836	0		Both	Yes No
SAN Load Source Specify	0837	0		Both	Yes No
#3676 Load Source Specify	0838	0		Both	Yes No
#3677 Load Source Specify	0839	0		Both	Yes No
#3678 Load Source Specify	0840	0		Both	Yes No
#4329 Load Source Specify	0841	0		Both	Yes No
Modem Cable -- US/Canada and General Use	1025	18		Both	Yes No
System Unique Identifier	1311	0		MES	Yes No
200V 16A 4.3m (14-Ft) TL Line Cord	1406	0		Support	Yes No
125V 4.3m (14-Ft) Line Cord	1413	0		Support	Yes No
200V 1.8m (6-Ft) Locking Line Cord	1414	0		Support	Yes No
200V 1.8m (6-Ft) Watertight Line Cord	1415	0		Support	Yes No
200V 4.3m (14-Ft) Locking Line Cord	1416	0		Support	Yes No
200V 4.3m (14-Ft) Watertight Line Cord	1417	0		Support	Yes No
4.3m 200V/16A Power Cord	1418	18		Support	Yes No

S. Africa

4. 3m 200V/16A Power Cord Israel	1419	18	Support	Yes	No
4. 3m 200V/16A Power Cord EU/Asia	1420	18	Support	Yes	No
4. 3m 200V/16A Power Cord CH/DK	1421	18	Support	Yes	No
200V 1. 8m (6-F) Locking Line Cord	1424	0	Support	Yes	No
200V 1. 8m (6-Ft) Watertight Line Cord	1425	0	Support	Yes	No
200V 4. 3m (14- Ft) Locking Line Cord	1426	0	Support	Yes	No
200V 4. 3m (14- Ft) Watertight Line Cord	1427	0	Support	Yes	No
4. 3m 200V/10A Power Cord EU/Asia	1439	18	Support	Yes	No
4. 3m 200V/10A Power Cord Denmark	1440	18	Support	Yes	No
4. 3m 200V/10A Power Cord S. Africa	1441	18	Support	Yes	No
4. 3m 200V/10A Power Cord Swiss	1442	18	Support	Yes	No
4. 3m 200V/10A Power Cord UK	1443	18	Support	Yes	No
4. 3m 200V/10A Power Cord Israel	1445	18	Support	Yes	No
4. 3m 200V/32A Power Cord EU 1-PH	1449	314	Support	Yes	No
4. 3m 200V/16A Power Cord EU 2-PH	1450	18	Support	Yes	No
200V (6- Ft) 1. 8m Line Cord	1451	52	Support	Yes	No
200V (14- Ft) 4. 3m Line Cord	1452	52	Support	Yes	No
200V (6- Ft) 1. 8m Locking Line Cord	1453	52	Support	Yes	No
200V 12A (14- Ft) 4. 3m TL Line Cord	1454	262	Support	Yes	No
200V (6- Ft) 1. 8m Watertight Line Cord	1455	262	Support	Yes	No
200V (14- Ft) 4. 3m Watertight Line Cord	1456	262	Support	Yes	No
200V (6- Ft) 1. 8m Upper Line Cord	1457	0	Support	Yes	No
200V (6- Ft) 1. 8m Upper Locking Cord	1458	0	Support	Yes	No
200V (6- Ft) 1. 8m Upper Locking Cord	1459	262	Support	Yes	No
3m Copper RIO Cable	1460	589	MES	Yes	No

6m Copper RIO Cable	1461	720	MES	Yes No
15m RIO Cable	1462	1,440	MES	Yes No
30m SPCN Cable	1466	0	Support	Yes No
6m RIO to RIO-2 Cable	1474	851	MES	Yes No
10m RIO to RIO-2 Cable	1475	1,572	MES	Yes No
4.3m 200V/12A Pwr Cd UK	1476	18	Support	Yes No
4.3m 200V/16A Pwr Cd	1477	262	Support	Yes No
Remote I/O Cable, 15M	1485	1,440	Both	Yes No
3m RIO to RIO-2 Cable	1487	540	Both	Yes No
IPCS Keyboard/Mouse for NT	1700	118	Support	Yes No
0.6 Meter 12X Cable	1829	463	Both	Yes No
1.5 Meter 12X cable	1830	529	Both	Yes No
8.0 Meter 12X Cable	1834	960	Both	Yes No
3.0 Meter 12X Cable	1840	629	Both	Yes No
Ultra 320 SCSI Cable 1 Meter	2124	166	Both	Yes No
Ultra 320 SCSI Cable 3 Meter	2125	185	Both	Yes No
Ultra 320 SCSI Cable 5 Meter	2126	206	Both	Yes No
Ultra 320 SCSI Cable 10 Meter	2127	278	Both	Yes No
Ultra 320 SCSI Cable 20 Meter	2128	437	Both	Yes No
0.55 Meter Ultra 320 SCSI Cable	2138	100	Both	Yes No
Primary OS -- IBM i Operating System	2145	0	Both	N/A No
Primary OS -- AIX	2146	0	Both	N/A No
Primary OS -- Linux	2147	0	Both	N/A No
PCI Ultra Mag Media Controller	2749	1,702	Support	Yes No
PCI-X Ultra RAID Disk Controller	2757	9,425	Support	Yes No
PCI-X Ultra4 RAID Disk Controller	2780	8,116	Support	Yes No
PCI-X Fibre Chan Disk Controller	2787	8,116	Support	Yes No
PCI IOP	2844	2,749	Both	Yes No
PCI IOP for SAN Load Source	2847	3,273	Both	Yes No
PCIe 2-Line WAN w/Modem	2893	766	Both	Yes No
PCIe 2-Line WAN w/Modem CIM	2894	766	Both	Yes No
SAS Cable (X) Adapter to SAS Enclosure, Dual Controller/Dual Path 3M	3661	199	Both	Yes No
SAS Cable (X) Adapter to SAS Enclosure, Dual	3662	399	Both	Yes No

Controller/Dual Path 6M					
SAS Cable (X) Adapter to SAS Enclosure, Dual Controller/Dual Path 15M	3663	699	Both	Yes	No
SAS Cable (YR) - 1M	3667	125	Both	Yes	
69. 7GB 15k rpm SAS Disk Drive	3676	659	Both	Yes	No
139. 5GB 15k rpm SAS Disk Drive	3677	1, 299	Both	Yes	No
283. 7GB 15k rpm SAS Disk Drive	3678	1, 999	Both	Yes	No
SAS Cable (AE) Adapter to Enclosure, single controller/single path 3M	3684	199	Both	Yes	No
SAS Cable (AE) Adapter to Enclosure, single controller/single path 6M	3685	399	Both	Yes	No
SAS Cable (Y0) Adapter to SAS Enclosure, Single Controller/Dual Path 15 M	3694	699	Both	Yes	No
External xSeries(R) Attach	3704	0	Support	N/A	No
PCI IOP	3705	2, 520	Support	Yes	No
DVD-ROM	3706	785	Support	Yes	No
30GB 1/4-Inch Cartridge Tape	3707	5, 236	Support	Yes	No
50GB 1/4-Inch Cartridge Tape	3708	7, 855	Support	Yes	No
PCI 100/10Mbps Ethernet IOA	3709	539	Support	Yes	No
35. 16GB 10k rpm Disk Unit	4319	982	Support	Yes	No
35. 16GB 15k rpm Disk Unit	4326	1, 570	Support	Yes	No
70. 56GB 15k rpm Disk Unit	4327	999	Both	Yes	No
141. 12GB 15k rpm Disk Unit	4328	1, 299	Both	Yes	No
282. 25GB 15k rpm Disk Unit	4329	2, 799	Both	Yes	No
DVD-RAM	4430	2, 095	Support	Yes	No
50GB 1/4-Inch Cartridge Tape	4487	7, 855	Support	Yes	No
DVD-RAM	4630	2, 095	Support	Yes	No
DVD-RAM	4633	2, 095	MES	Yes	No
PCI Twinaxial Workstn IOA	4746	982	Both	Yes	No
PCI Crypto Coprocessor	4801	6, 545	Support	Yes	No
PCI Crypto Accelerator	4805	7, 855	Support	Yes	No
PCI Integ xSeries Server	4812	3, 639	Support	Yes	No
PCI Integ xSeries Server	4813	3, 639	Support	Yes	No
CBU SPECIFY	4891	0	MES	N/A	No
Single 5250 Enterprise Enablement	4990	50, 000	Both	N/A	No
Full 5250 Enterprise Enablement	4991	\$200, 000	Both	N/A	No

Software Preload Required	5000	0		Initial	N/A	No
Customer Solution Center -- Rochester Mfg	5002	0		Initial	N/A	No
PCI-X Expansion Unit	5088	15,709	\$200	Support	Yes	No
PCI-X Expansion Tower	5094	20,815	428	Support	Yes	No
PCI-X Exp Tower (no disk)	5096	18,982	200	Support	Yes	No
30-Disk Expansion Feature	5108	11,782	456	Both	No	No
Dual Line Cords -- Tower	5115	3,927		Both	Yes	No
Dual Line Cords -- 5294 Tower	5116	3,927		Both	Yes	No
Redundant Power and Cooling	5138	1,767		Both	Yes	No
Power Dist Unit 1 Phase NEMA	5160	1,309		Support	Yes	No
Power Dist Unit 1 Phase IEC	5161	1,309		Support	Yes	No
Power Dist Unit 2 of 3 Phase	5162	1,309		Support	Yes	No
Power Dist Unit -- 3 Phase	5163	1,309		Support	Yes	No
1.8m I/O Tower	5294	69,251	\$1,767	Support	Yes	No
1.8m I/O Tower (no disk)	5296	42,022	399	Support	Yes	No
Utility Billing for FC# 7380 with IBM i -- 100 processor minutes	5480	1,130		MES	Yes	No
Utility Billing for FC# 5620 with IBM i -- 100 processor minutes	5481	950		MES	Yes	No
Utility Billing for FC# 5621/#5622 with IBM i -- 100 processor minutes	5482	1,030		MES	Yes	No
On/Off Processor Billing for FC#5620 with IBM i -- 1 processor day	5483	570		MES	Yes	No
On/Off Processor Billing for FC#5621/5622 with IBM i -- 1 processor day	5484	620		MES	Yes	No
On/Off Processor Billing for FC#7380 with IBM i -- 1 processor day	5485	680		MES	Yes	No
Sys Console on OP Console	5544	0		Both	Yes	No
Sys Console 100Mbps Ethernet	5548	0		Support	Yes	No
Sys Console On HMC	5550	0		Both	Yes	No
Sys Console-Ethernet No IOP	5553	0		Both	Yes	No
Mirror 35GB Disk/Controller Pkg	5554	18,623		Support	Yes	No
Mirror 70GB Disk/Controller Pkg	5555	23,810		Support	Yes	No
Mirror 141GB Disk/Controller Pkg	5556	36,377		Support	Yes	No
Mirror 35GB Drawer Package	5560	36,000	\$229	Support	Yes	No

Mirror 70GB Drawer Package	5561	44,378	229	Support	Yes	No
2780 Controller w/Aux Write Cache	5580	9,157		Support	Yes	No
2757 Controller w/Aux Write Cache	5581	10,466		Support	Yes	No
5777 Controller w/Aux Write Cache	5583	9,157		Both	Yes	No
2780 Controller w/Aux Write Cache	5590	9,157		Support	Yes	No
2757 Controller w/Aux Write Cache	5591	9,157		Support	Yes	No
PCI-X Ultra Tape Controller	5702	861		Support	Yes	No
PCI-X Fibre Channel Tape Controller	5704	3,464		Support	Yes	No
IBM Single Bus Ultra 320 SCSI Repeater Card	5741	660	\$98	Both	Yes	No
IBM Dual Bus Ultra 320 SCSI Repeater Card	5742	1,321	98	Both	Yes	No
4Gbps Fibre Channel (2-Port)	5749	3,308		Both	Yes	No
PCI-X Fibre Chan Disk Controller	5760	5,495		Both	Yes	No
PCI-X Fibre Chan Tape Controller	5761	3,464		Both	Yes	No
10 Gigabit Ethernet-LR PCI Express Adapter	5772	6,276		Both	Yes	No
PCI-X Disk Controller-90MB No IOP	5776	2,617		Both	Yes	No
PCI-X Disk Controller-1.5GB No IOP	5777	8,116		Both	Yes	No
PCI-X EXP24 Ctl-1.5GB No IOP	5778	11,127		Both	Yes	No
PCI-X EXP24 Ctl-1.5GB No IOP	5782	11,193		Both	Yes	No
Total Storage(R) EXP24 Disk Dwr	5786	7,200	\$62	Both	Yes	No
PCI Expansion Drawer	5790	5,354	86	Both	Yes	No
PCI-DDR 12X Expansion Drawer	5796	6,545	86	Both	Yes	No
PCI-X DDR Dual Channel Ultra320 SCSI Adapter	5806	768		Both	Yes	No
PCI-X DDR Dual -- x4 3Gb SAS RAID Adapter	5902	2,500		Both	Yes	No
Alternate SAS controller for 3 of 6 internal SAS Disk Slots	5909	1,400		Both	Yes	No
PCI-X DDR Dual -- x4 SAS Adapter	5912	1,092		Both	Yes	No
Opt Front Door for 1.8m Rack	6068	450		MES	Yes	No
Opt Front Door for 2.0m Rack	6069	550		MES	Yes	No
1.8m Rack Trim Kit	6246	299		MES	Yes	No
2.0m Rack Trim Kit	6247	299		MES	Yes	No

1. 8m Rack Acoustic Doors	6248	4, 650	MES	Yes No
2. 0m Rack Acoustic Doors	6249	4, 650	MES	Yes No
RI0-2 Bus Adapter	6417	1, 059	Support	Yes No
RI0-2 Remote I/O Loop Adapter for #5790	6438	1, 191	Both	Yes No
Dual - port 12X Channel Attach-Short Run	6446	762	Both	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6451	7	Support	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6455	0	Support	Yes No
Dual - port 12X Channel Attach-Long Run	6457	3, 309	Both	Yes No
3. 7m (12- Ft) 250V/10A RA Pwr Cd	6459	7	Both	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6461	7	Both	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6462	7	Support	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6463	7	Support	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6464	7	Support	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6465	7	Support	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6466	7	Support	Yes No
4. 3m (14- Ft) 250V/10A Power Cord	6467	7	Support	Yes No
4. 3m (14- Ft) 3PH/24A Power Cord	6489	275	MES	Yes No
4. 3m (14- Ft) 1PH/48A Pwr Cord	6491	400	MES	Yes No
4. 3m (14- Ft) 1PH/48- 60A Pwr Cord	6492	400	MES	Yes No
Optional Rack Security Kit	6580	180	MES	Yes No
Modem Tray for 19- Inch Rack	6586	250	MES	Yes No
4. 3m (14- Ft) 1PH/24- 30A Pwr Cord	6654	240	MES	Yes No
4. 3m (14- Ft) 1PH/24- 30A WR Pwr Cord	6655	449	MES	Yes No
4. 3m (14- Ft) 1PH/24A Power Cord	6656	240	MES	Yes No
4. 3m (14- Ft) 1PH/24A Power Cord	6657	250	MES	Yes No
RI0-2 Bus Adapter	6699	1, 047	Both	Yes No
PCI 2- Line WAN I/O No IOP	6805	556	Both	Yes No

PCI 4-Modem WAN IOA No IOP	6808	2, 095		Both	Yes No
PCI 2-Line WAN w/Modem No IOP	6833	766		Both	Yes No
Intelligent PDU+, 1 EIA Unit, Universal UTG0247 Connector	7109	1, 455		MES	Yes No
Power Distribution Unit	7188	1, 000		MES	Yes No
Quantity 150 of #2124	7204	24, 900		Both	Yes No
Quantity 150 of #2125	7205	27, 750		Both	Yes No
Quantity 150 of #2126	7206	30, 900		Both	Yes No
Quantity 150 of #2127	7207	41, 700		Both	Yes No
Quantity 150 of #2128	7208	65, 550		Both	Yes No
Quantity 150 of #2138	7213	15, 000		Both	Yes No
2GB CUoD Memory Activation	7272	3, 030		Both	Yes No
4GB CUoD Memory Activation	7273	6, 060		Both	Yes No
8GB CUoD Memory Activation	7274	12, 120		Both	Yes No
16GB CUoD Memory Activation	7275	24, 240		Both	Yes No
32GB CUoD Memory Activation	7276	48, 480		Both	Yes No
Dual I/O Unit Enclosure	7307	695		Both	Yes No
Dual I/O Unit Enclosure	7311	552		Support	Yes No
I/O Drawer Mounting Enclosure	7314	695		Both	Yes No
Quantity 150 of #4319	7504	147, 273		Support	Yes No
Quantity 150 of #4326	7508	235, 440		Support	Yes No
Quantity 150 of #4327	7509	149, 850		Both	Yes No
Quantity 150 of #4328	7510	194, 850		Both	Yes No
Quantity 150 of #4329	7511	419, 850		Both	Yes No
Quantity 150 of #5741	7514	99, 000	\$14, 678	Both	Yes No
Quantity 150 of #3676	7517	98, 850		Both	Yes No
Quantity 150 of #3677	7518	194, 850		Both	Yes No
Quantity 150 of #3678	7519	299, 850		Both	Yes No
2.0m Rack Side Attach Kit	7780	196		Support	Yes No
Side-by-Side for 1.8m Racks	7840	655		Support	Yes No
Ruggedize Rack Kit	7841	1, 964		Support	Yes No
Month Indicator	9461	NC		Initial	N/A No
Day Indicator	9462	NC		Initial	N/A No
Hour Indicator	9463	NC		Initial	N/A No
Minute Indicator	9464	NC		Initial	N/A No
Qty Indicator	9465	NC		Initial	N/A No
Countable Member Indicator	9466	NC		Initial	N/A No

Feature conversions

Feature conversions for 9406-570 to 9117-MMA administrative features

From FC:	To FC:	Parts returned	Purchase price
1641 -- 1.65 GHz Proc HW Upgr Feat	7380 -- 4.7 GHz Proc Card, 0/2 Core POWER6, 12 DDR2 Memory Slots	Yes	\$10,000
1651 -- 2.2 GHz Proc HW Upgr Feat	7380 -- 4.7 GHz Proc Card, 0/2 Core POWER6, 12 DDR2 Memory Slots	Yes	5,500

Feature conversions for 9406-570 to 9117-MMA capacity on demand features

From FC:	To FC:	Parts returned	Purchase price
7897 -- 570 CUoD Proc Activation	5403 -- One Processor Activation for Processor Feature #7380	No	\$17,000

Feature conversions for 9406-570 to 9117-MMA memory features

From FC:	To FC:	Parts returned	Purchase price
7663 -- 570 1GB Mem Activation	5680 -- Activation of 1GB DDR2 POWER6 Memory	No	\$ 758
7892 -- 2GB DDR2 Main Storage	7272 -- 2GB CUoD Memory Activation	Yes	1,515
4495 -- 4/8GB DDR2 Main Storage	7273 -- 4GB CUoD Memory Activation	Yes	3,030
7893 -- 4GB DDR2 Main Storage	7273 -- 4GB CUoD Memory Activation	Yes	3,030
4496 -- 8/16GB DDR2 Main Storage	7274 -- 8GB CUoD Memory Activation	Yes	6,060
7894 -- 8GB DDR2 Main Storage	7274 -- 8GB CUoD Memory Activation	Yes	6,060
4497 -- 16GB DDR2 Main Storage	7275 -- 16GB CUoD Memory Activation	Yes	12,120
4499 -- 16GB DDR2 Main Storage	7275 -- 16GB CUoD Memory Activation	Yes	12,120
4498 -- 32GB DDR2 Main Storage	7276 -- 32GB CUoD Memory Activation	Yes	24,240

Feature conversions for 9406-570 to 9117-MMA processor features

From FC:	To FC:	Parts returned	Purchase price
7260 -- 570 Enterprise Enablement	4990 -- Single 5250 Enterprise Enablement	No	\$ 0
7577 -- 570 Enterprise Enablement	4990 -- Single 5250 Enterprise Enablement	No	0
9286 -- Base Enterprise Enablement	4990 -- Single 5250 Enterprise Enablement	No	0
9299 -- Base 5250 Enterprise Enable	4990 -- Single 5250 Enterprise Enablement	No	0
7258 -- 570 Full Enterprise Enable	4991 -- Full 5250 Enterprise Enablement	No	0
7597 -- 570 Full Enterprise Enable	4991 -- Full 5250 Enterprise Enablement	No	0
9298 -- Full 5250 Enterprise Enable	4991 -- Full 5250 Enterprise Enablement	No	0
7618 -- 570 One Processor Activation	5403 -- One Processor Activation for Processor Feature #7380	No	12,500

Order now

To order, contact the Americas Call Centers, your local IBM representative, or your IBM Business Partner.

To identify your local IBM representative or IBM Business Partner, call 800-IBM-4YOU (426-4968).

Phone: 800-IBM-CALL (426-2255)
Fax: 800-2IBM-FAX (242-6329)
Internet: callserv@ca.ibm.com
Mail: IBM Teleweb Customer Support
ibm.com Sales Execution Center, Americas North
3500 Steeles Ave. East, Tower 3/4
Markham, Ontario
Canada
L3R 2Z1

Reference: YE001

The Americas Call Centers, our national direct marketing organization, can add your name to the mailing list for catalogs of IBM products.

Note: Shipments will begin after the planned availability date.

Trademarks

POWER, AIX 5L, POWER6, POWER5, POWER5+, and Electronic Service Agent are trademarks of International Business Machines Corporation in the United States or other countries or both.

AIX, xSeries, and TotalStorage are registered trademarks of International Business Machines Corporation in the United States or other countries or both.

Intel is a registered trademark of Intel Corporation.

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

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

This announcement is provided for your information only. For additional information, contact your IBM representative, call 800-IBM-4YOU, or visit the IBM home page at: <http://www.ibm.com>