

IBM Power System E950 server offers a unique blend of enterprise-class capabilities in a reliable, agile, space-efficient 4-socket 4U form factor that can deliver exceptional performance, and it is ideal for cloud deployments with built-in virtualization, flexible capacity, and high utilization

Table of contents

2 Overview	31 Publications
4 Key prerequisites	33 Technical information
4 Planned availability date	37 Terms and conditions
5 Description	44 Prices
23 Statement of general direction	44 Announcement countries
23 Product number	45 Corrections

At a glance

The IBM^(R) Power^(R) System E950 server (9040-MR9) is an enhanced 4-socket, 4U server that provides a significant enhancement to IBM Power SystemsTM 4-socket servers. It is ideal for cloud deployments with built-in virtualization, flexible capacity, and guaranteed utilization.

Summary of Power System E950 features:

- IBM POWER9TM processor modules with 8, 10, 11, or 12 POWER9 processor cores
 - 8-core processor modules provide up to 32 cores per system typical 3.6 - 3.8 GHz (max).
 - 10-core processor modules provide up to 40 cores per system typical 3.4 - 3.8 GHz (max).
 - 11-core processor modules provide up to 44 cores per system typical 3.2 - 3.8 GHz (max).
 - 12-core processor modules provide up to 48 cores per system typical 3.15 - 3.8 GHz (max).
- 128 GB to 16 TB high-performance industry standard double data rate (ISDDR) memory
 - A two-processor module configuration provides 64 ISDIMM slots.
 - A four-processor module configuration provides 128 ISDIMM slots and a server maximum of 16 TB.
- PCIe slots
 - System unit supports up to 10 PCIe Gen4 slots.
 - One PCIe Gen3 reserved for default Ethernet adapter.
 - PCIe I/O Expansion Drawers support up to 48 additional PCIe slots.
- Standard PCIe3 SAS controllers (#EJ0K) for internal disk bays
 - Eight SFF (2.5 in.) SAS bays for HDDs or SSDs
- Four NVMe U.2 bays
- Integrated USB controller provides two front and two rear USB ports. External DVD plugs into front USB port.

- Additional storage can be attached.
 - Up to 1,536 SFF-2 bays for HDDs and SSDs with 64 EXP24SX or EXP12SX I/O drawers
 - Fibre Channel and LAN adapters
- Elastic capacity on demand for both cores and memory
- Power to Cloud Rewards
- Active Memory™ Expansion and Active Memory Mirroring for Hypervisor
- Packaged in a standard 19-inch rack-mount 4U drawer
- Support for IBM AIX^(R) and Linux^(R) operating systems

Overview

IBM Power Systems for enterprise combines cloud management with leadership performance, scale, availability, and security to cloud enable the most data-intensive, mission-critical workloads.

The new Power System E950 server has built-in virtualization, flexible capacity, and high utilization, ideal for growing, medium-size businesses and as a departmental server or data center building block for large enterprises. Built with the new POWER9 architecture, the E950 provides improvements in application delivery and IT services through increased throughput, reduced response times, and increased memory and I/O bandwidths.

With its built-in virtualization and an ecosystem of open cloud solutions, you can cloud-enable any workload without migration and build a cloud that handles the most data-intensive workloads on earth on either a private, hybrid, or multicloud platform.

- Help reduce your operating costs and increase efficiency by consolidating your older, underutilized servers with the new compact, energy-efficient 4U form factor of the Power E950 server with its single backplane and single cooling domain.
- Help reduce maintenance and increase resource utilization and efficiency with up to 48 POWER9 cores and up to 16 TB DDR4 memory (4x more than IBM POWER8^(R) E850C).
- Help enable more agile IT that rapidly adjusts to the demands of an ever-changing business environment. IBM PowerVM^(R) supports dynamic response to workload changes and seamlessly provisions processor or memory capacity permanently, or only as long as required, with capacity on demand for processors and memory.

This new server generation has quadruple the memory footprint of POWER8 servers, making it an ideal platform for in-memory and data-centric applications. Changes in the memory subsystem and the use of industry standard memory DIMMs take POWER9 technology to the next level of balancing affordability and performance. Designed to run commercial, cognitive, and database workloads, POWER9 technology provides a highly competitive server platform. Client references indicate POWER^(R) servers help provide a robust and secure backbone for their IT infrastructure. More companies are using POWER technology in their IT infrastructure, from large data center down to shop level deployments.

By default, the Power Management mode of the Power E950 server is set to Max Performance. This mode can dynamically optimize the processor frequency at any given time based on CPU utilization and operating environmental conditions.

The IBM Power System E950 server offers Elastic Capacity on Demand (Elastic CoD) for both cores and memory, Active Memory Expansion, and Active Memory Mirroring for Hypervisor. The E950 server supports AIX and Linux operating systems and IBM PowerVM Enterprise. The Power System E950 server is a logical follow-on to the IBM Power System E850 and E850C servers.

Each Power System E950 server includes the following additional software and services:

- Cloud Management Console (CMC) for IBM Power Systems- CMC is a cloud-native platform that provides apps that give powerful insights into your Power Systems infrastructure across data centers and geographies. With no additional software or infrastructure setup, get single pane of glass views of your inventory, software levels, resource capacity utilization, as well as launch-in-context of your on-prem software such as PowerVC and PowerHA^(R). Each new Power E950 server purchase includes 12 months of CMC service, which can be renewed. For more information, see [IBM CloudTM Management Console](#) website.
- Open source cloud automation and configuration tooling for AIX. IBM has expanded its commitment to keep key open source cloud management packages updated and to provide timely security fixes to enable clients to leverage open source skills. Power System E950 server clients can take advantage of key packages recently provided to enable cloud automation, including:
 - **chef** automation for configuration, deployment, and management. IBM is collaborating with clients in this community to provide useful resources for using chef with AIX systems.
 - **yum** package management is available with repository access from both ftp and https protocols. rpm is also updated to enable automatic dependency discovery.
 - **cloud-init** and all dependencies are available in the repository and includes support for licensed AIX users. For more information, see the [AIX Toolbox for Linux Applications](#) website.
- Securely connect system-of-record workloads and data-to-cloud native applications. IBM API Connect^(R) and IBM WebSphere^(R) Connect provide secure connectivity to cloud-based applications, giving clients the ability to rapidly develop new applications and services, thereby helping accelerate their time to value. IBM's Power to Cloud services can help clients get started with these solutions and in designing new applications that leverage IBM Cloud. Cloud enables clients to rapidly build, deploy, and manage their cloud applications, while tapping a growing ecosystem of available services and runtime frameworks.
- IBM Power to Cloud Reward Program. The IBM Power to Cloud Reward Program helps clients design, build, and deliver cloud platforms on Power Systems. Clients purchasing an Power System E950 server earn 5,000 IBM Power to Cloud Reward Program points per system. The points can be used for a range of services to transform traditional IT platforms to private and hybrid cloud platforms. Power to Cloud Reward Program services offerings leverage the proven expertise of IBM Systems Lab Services consultants. For additional details, see the [IBM Power to Cloud Reward Program](#) website or contact your IBM representative, IBM Business Partner, or ibmsls@us.ibm.com.

Power System E950 hardware components

The Power System E950 server delivers a high-performing and reliable 4-socket system in the marketplace, with simultaneous multithreading of up to eight threads per core (SMT8). Each Power System E950 server's POWER9 single-chip module has two on-chip memory controllers and utilizes up to 128 MB off-chip eDRAM L4 cache to deliver 230 GB/sec of memory bandwidth per socket. I/O bandwidth is also increased to double that of the predecessor POWER8 systems with the introduction of PCIe Gen4 integrated controllers.

The new Power System E950 server utilizes 8-core, 10-core, 11-core, or 12-core POWER9 processor modules and delivers a system with up to 32 cores typical 3.6 - 3.8 GHz (max), up to 40 cores typical 3.4 - 3.8 GHz (max), up to 44 cores typical 3.2 - 3.8 GHz (max), or up to 48 cores typical 3.15 - 3.8 GHz (max). A minimum of two processor modules must be installed in each system, with a minimum quantity of one processor module's cores activated.

An Power System E950 server can utilize 8 GB, 16 GB, 32 GB, 64 GB, or 128 GB DDR4 ISDIMMs (up to 16 industry standard (IS) DIMMs per socket) to enable a

maximum configuration of 16 TB. A minimum of 50% or 128 GB, whichever is higher, of installed memory must be activated.

The PCIe capabilities provide up to eight PCIe Gen4 x16, two PCIe Gen4 x8, and one PCIe Gen3 x8. PCIe slots support concurrent maintenance (hot plug). Additional PCIe adapters can be attached through up to four PCIe I/O Expansion Drawers, providing up to 48 additional PCIe Gen3 slots.

Each Power System E950 server has eight SFF (2.5 in.) SAS bays for HDDs or SSDs, plus four NVMe 2.5-inch bays. One or two PCIe SAS adapters support these SAS bays. One of these SAS options must be selected:

- DASD Backplane with no HDD/SDD drive selected. No PCIe SAS adapter is required.
- Base DASD backplane together with one SAS PCIe adapter and select SAS drives.
- Split DASD backplane together with two SAS PCIe adapters and select SAS drives.

NVMe drive options can be selected with any of the above SAS options

All SAS bays support concurrent maintenance (hot plug). In addition to the eight internal drives, more HDD or SSD drives can be attached through EXP24SX or EXP12SX expansion drawers.

The external USB DVD plugs into one of the two front USB ports.

Every Power System E950 server comes standard with phase-redundant, power regulators for processors, memory, disk, and I/O. Hot-plug, redundant power supplies and a hot-plug Time of Day battery are also standard, and Active Memory Mirroring for Hypervisor is available as an option.

IBM Proactive Support for Power Systems

Clients have found significant value in IBM's Proactive Support offerings on mission-critical systems, as this provides personalized support, proactive recommendations, and accelerated response times versus standard support. As a result, IBM is including the IBM Proactive Support in the default configuration for all Mid-Range and Enterprise IBM POWER9 Systems - for IBM AIX, IBM i, Linux, and SAP HANA workloads. Other configurations are also available.

Key prerequisites

One of the following operating systems:

- AIX
- Linux: Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES)

For more information, see the [Software requirements](#) section.

Planned availability date

August 17, 2018, with the following limitations per system:

- EMX0 Initial/MES/Both max = 2
- EMXG Initial/MES/Both max = 4

November 16, 2018, for the following features per system:

- EMX0 Initial/MES/Both max = 4

Availability of programs with an encryption algorithm in France is subject to French government approval.

Description

Summary of features

The following features are available on the Power System E950 server:

- The Power System E950 server supports 16 to 48 processor cores with two or four POWER9 processor modules:
 - 8-core processor modules provide up to 32 cores per system typical 3.6 - 3.8 GHz (max).
 - 10-core processor modules provide up to 40 cores per system typical 3.4 - 3.8 GHz (max).
 - 11-core processor modules provide up to 44 cores per system typical 3.2 - 3.8 GHz (max).
 - 12-core processor modules provide up to 48 cores per system typical 3.15 - 3.8 GHz (max).
- 128 GB - 16 TB high-performance DDR4 memory with L4 cache:
 - 8 GB ISDIMM Memory (#EM6A)
 - 16 GB ISDIMM Memory (#EM6B)
 - 32 GB ISDIMM Memory (#EM6C)
 - 64 GB ISDIMM Memory (#EM6D)
 - 128 GB ISDIMM Memory (#EM6E)
 - Optional Active Memory Expansion (#EMAM)
- Choice of three storage backplane features. Each Power E950 server has eight SFF (2.5 in.) SAS bays for HDDs or SSDs, plus four NVMe 2.5-inch bays. One or two PCIe SAS adapter support these SAS bays. One of these SAS options must be selected:
 - DASD Backplane with no HDD/SDD drive selected. No PCIe SAS adapter is required (#EJOB).
 - Base DASD backplane together with one SAS PCIe adapter and select SAS drives (#EJBB).
 - Split DASD backplane together with two SAS PCIe adapters and select SAS drive (#EJSB).
- Up to 11 hot-swap PCIe slots in the system unit:
 - Four, six, or eight x16 Gen4 full-height, half-length slots.
 - Two x8 Gen4 full height, half-length slots.
 - One x8 Gen3 full-height, half-length slot generally is used for an Ethernet adapter.
 - With two processor modules, there are seven PCIe slots; with three modules, there are nine PCIe slots; and with four modules, there are eleven PCIe slots in the system unit.
- The PCIe I/O Expansion Drawer (#EMX0) expands the number of full-high, hot-swap slots:
 - Up to two PCIe drawers with two processor modules (maximum 31 slots on the server)
 - Up to four PCIe drawers with four processor modules (maximum 51 slots on the server)
- Up to 64 EXP24SX SFF (#ESLS) or EXP12SX LFF (#ESLL) Drawers can be attached, providing up to 1,536 SAS bays for disks or SSDs.
- System unit I/O (integrated I/O)
 - USB ports: four 3.0 (two front) for general use and 2.0 (rear) for limited use
 - HMC ports: two 1 GbE RJ45

- System (serial) port: one RJ45
- Four hot-plug and redundant power supplies 2000 W (200 - 240 V AC) (#EB3M)
- System unit only 4U in a 19-inch rack-mount hardware
- Primary operating systems:
 - AIX (#2146) (small-tier licensing)
 - Linux (#2147): RHEL, SLES, and Ubuntu

Processor modules

- The Power E950 server supports 16 to 48 processor cores:
 - 8-core typical 3.6 - 3.8 GHz (max) #EPWR POWER9 Processor
 - 10-core typical 3.4 - 3.8 GHz (max) #EPWS POWER9 Processor
 - 11-core typical 3.2 - 3.8 GHz (max) #EPWY POWER9 Processor
 - 12-core typical 3.15 - 3.8 GHz (max) #EPWT POWER9 Processor
- A minimum of two and a maximum of four processor modules are required for each system. The modules can be added to a system at a later time through an MES order, but will require scheduled downtime to install. All processor modules in one server must be the same gigahertz (same processor module feature number). They cannot be mixed.
- Permanent CoD processor core activations are required for the first processor module in the configuration and are optional for the second and fourth modules. Specifically:
 - Two or four 8-core, typical 3.6 - 3.8 GHz (max) processor modules (#EPWR) require eight processor core activations (#EPWV) at a minimum.
 - Two or four 10-core, typical 3.4 - 3.8 GHz (max) processor modules (#EPWS) require ten processor core activations (#EPWW) at a minimum.
 - Two or four 11-core, typical 3.2 - 3.8 GHz (max) processor modules (#EPWY) require eleven processor core activations (#EPN3) at a minimum.
 - Two or four 12-core, typical 3.15 - 3.8 GHz (max) processor modules (#EPWT) require twelve processor core activations (#EPWX) at a minimum.
- Temporary CoD capabilities are optionally used for processor cores that are not permanently activated:
 - 90 Days Elastic CoD Processor Core Enablement (#EP9T).
 - 1 and 100 Processor Day Elastic CoD billing for #EPWR (#EPN0, #EPN1).
 - 1 and 100 Processor Day Elastic CoD billing for #EPWS (#EPN5, #EPN6).
 - 1 and 100 Processor Day Elastic CoD billing for #EPWY (#EPN8, #EPN9).
 - 1 and 100 Processor Day Elastic CoD billing for #EPWT (#EPNK, #EPNL).
 - 100 Processor-minutes Utility CoD billing: for #EPWR (#EPN2), for #EPWS (#EPN7), for #EPWY (#EPNN) or for #EPWT (#EPNM).
 - An HMC is required for temporary CoD.

Notes:

- A minimum of two processor features is required on each system.
- Mixing of different processor features on the same system is *not* allowed.

System memory

- 128 GB - 16 TB high-performance 1600 MHz DDR4 ECC memory with L4 cache:
 - 8 GB ISDIMM Memory (#EM6A)
 - 16 GB ISDIMM Memory (#EM6B)
 - 32 GB ISDIMM Memory (#EM6C)
 - 64 GB ISDIMM Memory (#EM6D)
 - 128 GB ISDIMM Memory (#EM6E)

- As the client memory requirements increase, the system capabilities are increased as follows:
 - With two processor modules installed, 64 ISDIMM slots are available; minimum memory is 128 GB.
 - With four modules, 128 ISDIMM slots are available; minimum memory is 256 GB. Thirty two ISDIMMs are available per socket.
 - The more ISDIMM slots that are filled, the larger the available bandwidth available to the server.

Permanent CoD memory activations are required for at least 50% of the physically installed memory or 128 GB of activations, whichever is larger. Use 1 GB activation (#EMAP) and 100 GB activation (#EMAQ) features to order permanent memory activations.
- Temporary CoD for memory is available for memory capacity that is not permanently activated:
 - 90 Days Elastic CoD Memory Enablement (#EM9U).
 - 8 GB-Day billing for Elastic CoD memory (#EMJE).
 - 800 GB-Day billing for Elastic CoD memory (#EMJF).
 - An HMC is required.

Temporary CoD Memory Days can also be acquired through IBM Marketplace after system installation. For more information about new Elastic CoD features, see the [IBM Digital MarketPlace](#) website.

Notes:

- Memory is ordered in quantity of 8 of the same memory feature.
- Second memory riser card of each processor can have different ISDIMM capacity from that on first riser but it is recommended that the second riser card has the same type and size DIMMs as the first riser card. It is also recommended that the amount of memory per processor should be the same or about the same.
- A minimum of one memory riser with eight ISDIMMs for each installed POWER9 processor is required.
- The minimum memory supported per two POWER9 processors installed is 128 GB.
- The minimum memory supported per four POWER9 processors installed is 256 GB.

I/O support

PCIe slots

The Power System E950 server has up to 10 PCIe Gen4 and 1 PCIOe Gen3 general purpose hot-plug slots, providing configuration flexibility and expandability. Eight adapter slots are x16 Gen4, two adapter slots are x8 Gen4, and one adapter slot x8 Gen3 (generally reserved for base Ethernet adapter). All adapter slots are full height, half length.

The number of slots supported vary by the number of processor modules.

Processor modules	2 socket	4 socket
x16 Gen4 slots (CAPI capable)	4	8
x8 Gen4 slots	2	2
x8 Gen3 slots	1	1

Notes:

- The PCIe Gen3 slot, C6, is reserved for an Ethernet adapter to help ensure proper manufacture and test of the server.
- Blind-swap cassettes (BSC) are used for adapters in the system unit.
- All PCIe slots in the system unit are SR-IOV capable.

- PCIe slot C12 must contain the SAS RAID adapter (#EJ0K) when controlling the Base internal SAS drive bays. The SAS RAID adapter in this slot will not support a Dual Storage Adapter configuration.
- PCIe slots C12 and C9 must each contain a SAS RAID controller (#EJ0K) when controlling the Split internal SAS drive bays. The SAS RAID adapters in these slots will not support a Dual Storage Adapter configuration.

The x16 slots can provide up to twice the bandwidth of x8 slots because they offer twice as many PCIe lanes. PCIe Gen4 slots can support up to twice the bandwidth of a PCIe Gen3 slot and up to four times the bandwidth of a PCIe Gen2 slot, assuming an equivalent number of PCIe lanes. PCIe Gen1, PCIe Gen2, PCIe Gen3, and PCIe Gen4 adapters can be plugged into a PCIe Gen4 slot, if that adapter is supported. The x16 slots can be used to attach PCI Gen3 I/O Expansion drawers.

This server is smarter than earlier servers about energy efficiency for cooling the PCIe adapter environment. They can sense which IBM PCIe adapters are installed in their PCIe slots. If an adapter is known to require higher levels of cooling, they automatically speed up fans to increase airflow across the PCIe adapters.

SAS bays and storage backplane options

Clients have a choice of three storage features. Each Power E950 server has eight SFF (2.5 in.) SAS bays for HDDs or SSDs, plus four NVMe 2.5-inch bays. One or two PCIe SAS adapter support these SAS bays. One of these SAS options must be selected:

- DASD Backplane with no HDD/SDD drive selected. No PCIe SAS adapter is required (#EJ0B)
- Base DASD backplane together with one SAS PCIe adapter and select SAS drives (#EJBB)
- Split DASD backplane together with two SAS PCIe adapters and select SAS drive (#EJSB)

NVMe drive options can be selected with any of the above SAS options.

The base and the split backplane options provides SFF-3 SAS bays in the system unit. These 2.5-inch or small form factor (SFF) SAS bays can contain SAS drives (HDD or SSD) mounted on a tray or carrier. Thus, any drives that are designated for SFF-1, or SFF-2 bays do not fit in an SFF-3 bay. All SFF-3 bays support concurrent maintenance or hot-plug capability. These backplane options support HDDs or SSDs or a mixture of HDDs and SSDs in the SFF-3 bays. If mixing HDDs and SSDs, they must be in separate arrays.

The base and the split options can offer different drive protection options: RAID 0, RAID 5, RAID 6, or RAID 10. RAID 5 requires a minimum of three drives of the same capacity. RAID 6 requires a minimum of four drives of the same capacity. RAID 10 requires a minimum of two drives. Hot spare capability is supported with RAID 5, RAID 6, or RAID 10.

Note:

- All three of these backplane options are supported by AIX, Linux, and VIOS. It is highly recommended that the drives be protected, but not required. If the client needs a change after the server is already installed, the backplane option can be changed. Scheduled downtime is required to remove the existing storage backplane and install a different backplane.

USB DVD external and boot devices

Power System 9040-MR9 provides two front USB ports, mainly to support the external USB DVD.

System boot is supported through these options:

1. NVMe drives

2. Internal SAS drives
3. Disk or SSD located in an EXP24SX or EXP12SX drawer attached to a PCIe SAS adapter
4. A network through LAN adapters
5. A SAN attached to Fibre Channel or FCoE adapters and indicated to the server by the 0837 specify code
6. USB front port for external USB based DVD
7. USB front port for USB memory key / flash drive

Assuming option 1 above, the minimum system configuration requires at least one NVMe disk drive in the system for AIX and Linux. If you are not using option 1 or 2 above, an internal disk is not required.

I/O drawer attachment

Clients migrating from earlier generation servers may have been using I/O drawers such as the GX++ attached feature 5802 or 5877 PCIe 12X I/O Drawers with PCIe Gen1 slots. Though most PCIe adapters in the feature 5802 or 5877 drawers can be moved to this server and its disk drives converted and moved to the feature ESLS EXP24SX drawer, the feature 5802 and 5877 drawers are not supported on this newer Power Systems technology-based server. Similarly, the GX++ attached EXP30 Ultra SSD Drawer (#EDR1 or #5888) is not supported.

Up to 64 EXP24SX or EXP12SX storage drawers can be attached to a Power System E950 server. With 24 drives per drawer, this means up to 1,536 SFF-2 drives are supported. A maximum of 16 EXP24SX per PCIe Gen3 I/O drawer is supported due to cable management considerations.

The older 3.5-inch-based EXP12S SAS Disk Drawer (#5886) and EXP24 SCSI Disk Drawer (#5786) are not supported.

IBM offers a 1U multimedia drawer that can hold one or more DVDs, tape drive, or RDX docking stations. The 7226-1U3 is the most current offering. The earlier 7216-1U2 and 7214-1U2 are also supported. Up to six of these multimedia drawers can be attached.

PCIe I/O Expansion Drawer

PCIe I/O Expansion Drawers (#EMX0) can be attached to the system unit to expand the number of full-high, hot-swap slots available to the server. The maximum number of PCIe I/O drawers depends on the number of processor modules physically installed. The maximum is independent of the number of processor core activations.

- Up to two PCIe drawers with two processor modules.
- Up to four PCIe drawers with four processor modules.
- The 19-inch 4 EIA (4U) PCIe I/O Expansion Drawer (#EMX0) and two PCIe fanout modules (#EMXG) provide 12 PCIe I/O full-length, full-height slots. One fanout module provides six PCIe slots labeled C1 through C6. Slots C1 and C4 are x16 slots, and C2, C3, C5, and C6 are x8 slots. Slots C1 and C4 of the 6-slot fanout module in a PCIe Gen3 I/O drawer are SR-IOV enabled.
- An EMX0 drawer can be configured with one or two EMXG fanout modules. Adding a second fanout module is not a hot-plug operation and requires scheduled downtime.
- PCIe Gen1, Gen2, and Gen3 full-high adapter cards are supported. The set of full-high PCIe adapters that are supported is found in the Sales Manual, identified by feature number. See the PCI Adapter Placement manual for the 9040-MR9 for details and rules associated with specific adapters supported and their supported placement in x8 or x16 slots.
- Each fanout module requires one PCIe Cable Adapter (#EJ08), which is placed in a x16 PCIe slot of the system unit.
- A cable pair attaches the PCIe Cable Adapter (#EJ08) to the fanout module. Feature ECC7 provides a pair of 3-meter optical cables with transceivers. Feature ECC8 provides a pair of 10-meter optical cables with transceivers. Feature ECCS

provides a pair of 3-meter copper cables. Two cables of identical length, or one feature, is required for each fanout module. Optical cables are smaller diameter and more flexible and can be longer than the copper cables. Copper cables are lower cost. Copper and optical cables have the same performance and reliability characteristics.

- The cable adapter (#EJ08) can be placed in any of the system units' available x16 slots and attached to any fanout module in any of the server's PCIe I/O drawers (#EMX0). The PCIe I/O drawer cannot be shared by two different servers.
- **Recommendation:** Locate any attached PCIe I/O Expansion Drawer in the same rack as the POWER9 server for ease of service. But the drawers can be installed in separate racks if the application or other rack content requires it. It is recommended to use 3-meter cables for PCIe drawers in the same rack as the system unit and 10-meter cables for drawers located in a different rack.
- Concurrent repair and add/removal of PCIe adapter cards is done by HMC-guided menus or by operating system support utilities.
- When the operating system is Linux, PowerVM is required for support of the I/O Expansion Drawer.

Rack cable brackets

The cable bracket is used to secure the communications cables so that they can be raised to remove the blind swap cassettes containing the PCIe adapters. To enable the server to be pulled forward on its rails for service access to memory, processors, and voltage regulator modules requires the disconnecting the cables from the server.

Integrated I/O ports

There are two HMC ports, one system port, and four USB ports (two general purpose and two limited use).

The two HMC ports are located on the service processor in the rear of the server and are RJ45. These ports support 1 Gb Ethernet connections leveraging the second-generation service processor technology called an *FSP2* in some IBM publications.

The one integrated system or serial port has an RJ45 style connector and is supported by AIX and Linux for attaching serial devices such as an asynchronous device like a console. If the device does not have an RJ45 connection, a converter cable such as feature 3930 can provide a 9-pin, D-shell connection. Note that serial devices can have very individual characteristics (different pin outs), and the feature 3930 may not be appropriate for all possible devices. In this case, the user should acquire an OEM converter cable appropriate for their device.

Integrated system port is not supported under AIX or Linux when the HMC ports are connected to an HMC. Either the HMC ports or the integrated system ports can be used, but not both. The integrated system port is supported for modem and TTY terminal connections by AIX or Linux. Any other application using serial ports requires a serial port adapter to be installed in a PCI slot. The integrated system port does not support HACMP™ configurations.

Two USB-3 ports are located on the front of the server for general client use, and two USB-3 ports on the rear are for limited client use. A converter cable, feature ECCF, provides a USB-to-9-pin D-Shell connection to some UPS. The IBM 9910 Uninterruptible Power Systems, models E66 and E67, are supported as purchase options with the 9040-MR9.

Racks

The IBM Power System E950 server is designed to fit a standard 19-inch rack. IBM Development has tested and certified the system for the following rack support:

- 7965-S42 (42U) (Field integration only)
- 7014-T42 (42U) (Field integration only)

Clients can choose to place the server in other racks if they are confident those racks have the strength, rigidity, depth, and hole-pattern characteristics that are needed. Clients should work with IBM Service to determine the appropriateness of other racks. The Power System E950 rails can adjust their depth to fit a rack that is 22.75 inches - 30.5 inches in depth based on the adjustable outer brackets going from the front rail to the rear rail. Racks from IBM include:

- 7014-T00 (36U) (Support only, cannot be ordered from IBM)

IBM Manufacturing does not support the use of other racks with the Power E950 initial system order or model upgrade. The 7965-S42 (#ECR0) and 7014-T42 (#0553) are 2-meter enterprise racks that provide 42U or 42 EIA of space.

Additional E950 PCIe I/O drawers (#EMX0) for an already installed server can be MES ordered with or without a rack.

Typical good cable management practices apply. For example, it is usually a good practice to leave 1U or 2U open at the top or bottom of the rack for cables exiting the rack or to hold extra cable lengths. Or if you have more than four PCIe I/O drawers in the same rack, carefully consider how many 4-port PCIe adapters have cables attached to all their ports. If you are not using something like the 8-inch extension for a 7965-S42 rack for extra space to lay the cables, service access can be a challenge, especially with thicker I/O cables. Or another example, if there are a lot of I/O cables in the rack, horizontally mounting the power distribution units (PDUs) versus mounting them vertically in the side pockets eases service access.

Rack front door options supported are the acoustic door (#6248 or #6249), the attractive geometrically accented door (#ERG7), and the cost-effective plain front door (#6068 or #6069). The front trim kit is also supported (#6263 or #6272).

Power distribution units

Using previously provided IBM PDU features 7188, 7109, and 7196 reduces the number of Power System E950 servers and other equipment that can be held most efficiently in a rack. The high-function PDUs provide more electrical power per PDU and thus offer better "PDU footprint" efficiency. In addition, they are intelligent PDUs that provide insight to actual power usage by receptacle and also provide remote power on/off capability for easier support by individual receptacle. The new PDUs are features EPTJ, EPTL, EPTN, and EPTQ.

IBM Manufacturing will only integrate the newer PDUs with the Power System E950 server. IBM Manufacturing does not support integrating earlier PDUs such as the feature 7188, 7109, or 7196. Clients can choose to use older IBM PDUs in their racks, but will have to install those earlier PDUs at their site.

	1-phase or 3-phase wye depending on country wiring standards	3-phase 208 V depending on country wiring standards
Nine C19 receptacles	EPTJ	EPTL
Twelve C13 receptacles	EPTN	EPTQ

- High Function 9xC19 PDU: Switched, Monitoring (#EPTJ). This is an intelligent, switched 200 - 240 V AC PDU with nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, making the nine C19 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Depending on country wiring standards, the PDU is single-phase or three-phase wye. The PDU can be mounted vertically in rack side pockets, or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. Device power cords with a C20 plug connect to C19 PDU receptacles and are ordered separately. One country-specific wall line cord is also ordered separately and attaches to a UTG524-7 connector on the front of the PDU. Supported line cords include features 6489, 6491, 6492, 6653, 6654, 6655, 6656, 6657, 6658, and 6667.
- High Function 9xC19 PDU 3-Phase: Switched, Monitoring (#EPTL). This is an intelligent, switched 208 V 3-phase AC PDU with nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, making the nine

C19 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. The PDU can be mounted vertically in rack side pockets, or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. Device power cords with a C20 plug connect to C19 PDU receptacles and are ordered separately. One wall line cord is provided with the PDU (no separate feature number) and has an IEC60309 60A plug (3P+G). The PDU supports up to 48 amps. Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password, and IBM strongly urges clients to change it upon installation.

- High Function 12xC13 PDU: Switched, Monitoring (#EPTN). This is an intelligent, switched 200 - 240 V AC PDU with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, making the twelve C13 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Depending on country wiring standards, the PDU is single-phase or three-phase wye. The PDU can be mounted vertically in rack side pockets, or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. Device power cords with a C14 plug connect to C13 PDU receptacles and are ordered separately. One country-specific wall line cord is also ordered separately and attaches to a UTG524-7 connector on the front of the PDU. Supported line cords include features 6489, 6491, 6492, 6653, 6654, 6655, 6656, 6657, 6658, and 6667. Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password, and IBM strongly urges clients to change it upon installation.
- High Function 12xC13 PDU 3-Phase: Switched, Monitoring (#EPTQ). This is an intelligent, switched 208 V 3-phase AC PDU with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack, making the twelve C13 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. The PDU can be mounted vertically in rack side pockets, or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. Device power cords with a C14 plug connect to C13 PDU receptacles and are ordered separately. One wall line cord is provided with the PDU (no separate feature number) and has a IEC60309 60A plug (3P+G). The PDU supports up to 48 amps. Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password, and IBM strongly urges clients to change it upon installation.

Lifting considerations

Three to four service personnel are required to manually remove or insert a system unit into a rack, given its dimensions and weight and content. To avoid the need for this many people to assemble at a client site for a service action, a lift tool can be very useful. Similarly, if the client has chosen to install this CSU (customer setup) system, similar lifting considerations apply.

The Power System E950 server has a maximum weight of 70.3 kg (155 lb). However, by temporarily removing the power supplies, fans, and RAID assembly, the weight is easily reduced to a maximum of 55 kg (121 lb).

When lowering the Power System E950 server onto its rails in the rack, the server must be tilted on one end about 15 degrees so that the pins on the server enclosure fit onto the rails. This equates to lifting one end of the server about 4 cm (1.6 in.). This can be done using a tip plate on a lift tool or manually adjusting the load on a lift tool or tilting during the manual lift. Consider the optional feature EB2Z Lift Tool.

PowerVM Enterprise Edition License Entitlement is now included with each Power System E950 server.

PowerVM Enterprise Edition is available as a hardware feature (#EPVV) and supports up to 20 partitions per core, VIOS, and multiple shared processor pools. It also offers Live Partition Mobility, Active Memory Sharing, and PowerVP™ performance monitoring.

Other PowerVM technologies include:

- System Planning Tool simplifies the process of planning and deploying Power Systems LPARs and virtual I/O.
- Virtual I/O Server (VIOS) is a single-function appliance that resides in an Power System partition. It facilitates the sharing of physical I/O resources between AIX, Linux, and client partitions within the server. VIOS provides shared Ethernet adapter (SEA) virtual I/O to client LPARs.
- Virtual SCSI (VSCSI) enables the sharing of physical storage adapters (Fibre Channel) and storage devices (disk and optical) between logical partitions.
- With virtual networking, a SEA enables connectivity between internal and external virtual LANs (VLANs); virtual Ethernet provides high-speed connections between partitions.

With Live Partition Mobility, you can move a running AIX or Linux VMs from one physical server to another with no downtime. Use this capability to:

- Evacuate workloads from a system before performing scheduled maintenance.
- Move workloads across a pool of different physical resources as business needs shift.
- Move workloads away from underutilized machines so that they can be powered off to save on energy and cooling costs.
- Move your workloads to new POWER8 systems for upgrades without having to schedule an outage. With Active Memory Sharing, memory is dynamically moved between running partitions for optimal resource usage.

Active Memory Expansion

Active Memory Expansion is an innovative technology supporting the AIX operating system that helps enable the effective maximum memory capacity to be larger than the true physical memory maximum. Compression/decompression of memory content can enable memory expansion up to 100% or more. This can enable a partition to do significantly more work or support more users with the same physical amount of memory. Similarly, it can enable a server to run more partitions and do more work for the same physical amount of memory.

Active Memory Expansion uses CPU resource to compress/decompress the memory contents. The trade-off of memory capacity for processor cycles can be an excellent choice, but the degree of expansion varies on how compressible the memory content is. It also depends on having adequate spare CPU capacity available for this compression/decompression.

POWER9 chips include a hardware accelerator designed to boost Active Memory Expansion efficiency and use less POWER core resource. The POWER9 accelerator includes some minor enhancements and also leverages POWER9 higher bandwidth and lower latency characteristics.

You have a great deal of control over Active Memory Expansion usage. Each individual AIX partition can turn on or turn off Active Memory Expansion. Control parameters set the amount of expansion desired in each partition to help control the amount of CPU used by the Active Memory Expansion function. An IPL is required for the specific partition that is turning on memory expansion. When turned on, monitoring capabilities are available in standard AIX performance tools such as lparstat, vmstat, topas, and svmon.

A planning tool is included with AIX, enabling you to sample actual workloads and estimate both how expandable the partition's memory is and how much CPU resource is needed. Any Power Systems model can run the planning tool. In addition, a one-time, 60-day trial of Active Memory Expansion is available to enable more exact memory expansion and CPU measurements. You can request the trial using the [Power Systems Capacity on Demand](#) web page.

Active Memory Expansion is enabled by chargeable hardware feature EMAM, which can be ordered with the initial order of the system node or as an MES order. A

software key is provided when the enablement feature is ordered, which is applied to the system node. An IPL is not required to enable the system node. The key is specific to an individual system node and is permanent. It cannot be moved to a different server.

The additional CPU resource used to expand memory is part of the CPU resource assigned to the AIX partition running Active Memory Expansion. Normal licensing requirements apply.

Capacity on demand

Several types of capacity on demand (CoD) processors are optionally available for the Power System E950 server. They help meet changing resource requirements in an on demand environment by using resources installed on the system but not activated.

Capacity upgrade on demand (CUoD) enables you to purchase additional permanent processor or memory capacity and dynamically activate it when needed.

There are two ways clients can order Elastic capacity on demand (Elastic CoD). With IBM Marketplace, clients can now order and provision Elastic CoD in minutes through the web. For more information about new Elastic CoD features, see the [IBM Digital Marketplace](#) website.

An additional way to order Elastic COD is through e-config features

Elastic 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.

With this offering, system administrators have 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 (8 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.

Use the following features to order activation features and support billing charges on the Power E950 (9040-MR9) server:

Processor module feature	Elastic CoD processor enablement feature	Elastic CoD AIX/Linux processor billing feature (1 and 100 proc-day)
EPWR (8-core module)	EP9T	EPN0, EPN1
EPWS (10-core module)	EP9T	EPN5, EPN6
EPWY (11-core module)	EP9T	EPN8, EPN9
EPWT (12-core module)	EP9T	EPNK, EPNL
Memory feature	Elastic CoD memory enablement feature	Elastic CoD memory billing feature (8 and 800 GB-day)
EM8P (16 GB)	EM9U	EMJE, EMJF
EM8Q (32 GB)	EM9U	EMJE, EMJF
EM8R (64 GB)	EM9U	EMJE, EMJF
EM8S (128 GB)	EM9U	EMJE, EMJF

Note: After CoD contracts are in place, inactive processor cores and inactive memory can be enabled for Elastic CoD usage.

The Elastic CoD processor consists of three steps: enablement, activation, and billing

- Elastic CoD enablement: Description

Before requesting temporary capacity on a server, you must enable it for Elastic CoD. To do this, order a no-charge enablement feature (MES only) and sign the required contracts or click to accept the terms on the CoD website. IBM will generate an enablement code, mail it to you, and post it on the web for you to retrieve and enter on your server. A processor enablement code lets you request up to 90 processor days of temporary unused CoD processor capacity for all your processor cores that have not been permanently activated. For example, if you have 16 processor cores that are not permanently activated, the processor enablement code allows up to 1,440 processor days (16 x 90). If you have reached or are about to reach the limit of 90 processor days per nonactivated processor core, obtain another enablement code through the web or by placing an order for another processor enablement code to reset the number of days you can request. Similarly, a memory enablement code lets you request up to 90 days of temporary unused CoD memory capacity for all your gigabytes of memory that have not been permanently activated. For example, if you had 100 GB of memory that was not permanently activated, the memory enablement code allows up to 9000 GB memory days (100 x 90). If you have reached the limit of 90 memory days per nonactivated memory, obtain another memory enablement code to reset the number of days you can request. Note that one 90-day enablement can be used over many months (more than three months) assuming your usage is modest.

See the [Power Systems Capacity on Demand](#) website for additional detail.

- Elastic activation requests:

When Elastic CoD temporary capacity is needed, simply use the HMC menu for Elastic CoD and specify how many of the inactive processors or how many gigabytes of memory you would like temporarily activated for some number of days. You will be billed for the days requested, whether the capacity is assigned to partitions or left in the shared processor pool. At the end of the temporary period (days you requested), you must ensure the temporarily activated capacity is available to be reclaimed by the server (not assigned to partitions) or you will be billed for any processor days not returned (per the contract you signed).

- Elastic CoD billing:

When the Elastic CoD is enabled, you are required to report billing data at least once a month (whether there is activity or not). This data is used to determine the proper amount to bill at the end of each billing period (calendar quarter). Failure to report billing data for use of temporary processor or memory capacity during a billing quarter will result in default billing equivalent to 90 processor days of temporary capacity. The sales channel will be notified of client requests for temporary capacity. As a result, the sales channel must order a quantity of billing features (using the appropriate billing features for each billable processor and memory day reported less any outstanding credit balance of processor and memory days).

For more information regarding registration, enablement, and usage of Elastic CoD, see the [Power Systems Capacity on Demand](#) website.

Note: Previously, Elastic CoD was called On/Off CoD. Some documentation still refers to this capability as "On/Off."

Utility CoD

Utility CoD provides 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 system node's shared processor pool, which then becomes available to the pool's resource manager. When the system node 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 enable the additional Utility CoD processors to be used. The system 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 website, and payment is based on reported usage. Utility CoD requires PowerVM Enterprise Edition to be active on the 9040-MR9.

Processor module feature	Utility CoD billing feature 100 processor minutes
EPWR (8-core module)	EPN2
EPWS (10-core module)	EPN7
EPWY (11-core module)	EPNN
EPWT (12-core module)	EPNM

For more information regarding registration, enablement, and use of Utility CoD, go to the [Utility Capacity on Demand](#) website.

Trial capacity on demand (Trial CoD): You can request either a standard or an exception trial at the [Trial Capacity on Demand](#) website.

Software licensing

For software licensing considerations with the various CoD offerings, see the latest revision of the *Capacity on Demand Planning Guide* at the [Power Systems Capacity on Demand](#) website.

IBM Power System E950 Solution Edition for Healthcare

The IBM Power System E950 Solution Edition for Healthcare provides a cost-effective 48-core/512 GB processor and memory activation feature package for eligible healthcare industry clients running approved ISV applications; for example, Epic.

- Solution Edition for Healthcare 12-core typical 3.15 - 3.8 GHz processor (#EHC4)
- Base Processor Activations (12) for #EHC4 (#ELAN)
- Base Memory Activation (128 GB) for #EHC4 (#EMAN)

The Power E950 Solution Edition for Healthcare minimum requirement is a server with 4x typical 3.15 - 3.8 GHz processor modules, 48 cores (all active), and 512 GB memory (all active).

Note: Additional hardware components can be added as desired following normal supported configuration rules.

For eligibility rules and registration of the Power Solution Edition for Healthcare by the sales channel, go to the [IBM Power Solution Editions for healthcare](#) website.

Reliability, Availability, and Serviceability

Reliability

The reliability of systems starts with components, devices, and subsystems that are designed to be highly reliable. 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 ensure product quality.

Power System E950 System RAS

Every Power System E950 server comes standard with phase-redundant power regulators for processors and memory. The memory regulators supply power to the DIMM memory riser cards and several processor voltage domains. Hot-plug, redundant power supplies and a concurrent maintenance of real time clock battery are also standard, and Active Memory Mirroring for Hypervisor is available as an option.

Processor VRM and Memory VRM are pluggable. Standby VRM and IO VRM also come with phase redundancy and are pluggable.

Power E950 comes with dual line cord redundancy along with n+1 power supply redundancy, also has n+1 fan rotor redundancy.

Memory subsystem RAS

The memory has error detection and correction circuitry designed such that the failure of any one specific DRAM module within an ECC word by itself can be corrected absent any other fault. The ECC word spans two DIMMs. In addition, there is a spare DRAM module per rank per DIMM pair.

Mutual surveillance

The service processor monitors the operation of the firmware during the boot process and also monitors the hypervisor for termination. The hypervisor monitors the service processor and reports a service reference code when it detects surveillance loss. In the PowerVM environment, it will perform a reset/reload if it detects the loss of the service processor.

Environmental monitoring functions

The Power Systems family does ambient and over temperature monitoring and reporting.

POWER9 processor functions

As in POWER8, the POWER9 processor has the ability to do processor instruction retry for some transient errors as well as provide core-contained checkstops for certain solid faults.

Cache availability

The L2 and L3 caches in the POWER9 processor in the memory buffer chip are protected with double-bit detect, single-bit correct error detection code (ECC). In addition, a threshold of correctable errors detected on cache lines can result in the data in the cache lines being purged and the cache lines removed from further operation without requiring a reboot in the PowerVM environment.

Modified data would be handled through Special Uncorrectable Error handling. L1 data and instruction caches also have a retry capability for intermittent errors and a cache set delete mechanism for handling solid failures.

Special Uncorrectable Error handling

Special Uncorrectable Error (SUE) handling prevents an uncorrectable error in memory or modified cache data 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 check stop. When and if data is used, I/O adapters controlled by an I/O hub controller would freeze if data were transferred to an I/O device, otherwise termination may be limited to the program/kernel or the if data is not owned by the hypervisor;

Uncorrectable error recovery

When the auto-restart option is enabled, the system can automatically restart following an unrecoverable software error, hardware failure, or environmentally induced (AC power) failure.

Serviceability

The purpose of serviceability is to efficiently repair the system while attempting to minimize or eliminate impact to system operation. Serviceability includes system installation, MES (system upgrades/downgrades), and system maintenance/repair. Depending upon the system and warranty contract, service may be performed by the client, an IBM representative, or an authorized warranty service provider.

The serviceability features delivered in this system help provide a highly efficient service environment by incorporating the following attributes:

- Design for Customer Set Up (CSU), Customer Installed Features (CIF), and Customer Replaceable Units (CRU)
- Error Detection and Fault Isolation (ED/FI)
- First Failure Data Capture (FFDC)
- Lightpath service indicators
- Service labels and service diagrams available on the system and delivered through IBM Knowledge Center
- Step-by-step service procedures documented in IBM Knowledge Center or available through the Hardware Management Console
- Automatic reporting of serviceable events to IBM through the Electronic Service Agent™ Call Home application
- CRU videos planned to be available on the web at general availability
- Mobile access to important customer service functions available by scanning a QR label

Service environment

In the PowerVM environment, the HMC is a dedicated server that provides functions for configuring and managing servers for either partitioned or full-system partition using a GUI or command-line interface (CLI) or REST API. An HMC attached to the system enables support personnel (with client authorization) to remotely, or locally to the physical HMC that is in proximity of the server being serviced, log in to review error logs and perform remote maintenance if required.

The POWER9 processor-based platforms support several service environments:

- Attachment to one or more HMCs or vHMCs is a supported option by the system with PowerVM. This is the default configuration for servers supporting logical partitions with dedicated or virtual I/O. In this case, all servers have at least one logical partition.
- For non-HMC systems.
 - Full-system partition with PowerVM: A single partition owns all the server resources and only one operating system may be installed. The primary service interface is through the operating system and the service processor.

Service interface

Support personnel can use the service interface to communicate with the service support applications in a server using an operator console, a graphical user interface on the management console or service processor, or an operating system terminal. The service interface helps to deliver a clear, concise view of available service applications, helping the support team to manage system resources and service information in an efficient and effective way. Applications available through the service interface are carefully configured and placed to give service providers access to important service functions.

Different service interfaces are used, depending on the state of the system, hypervisor, and operating environment. The primary service interfaces are:

- LEDs
- Operator Panel
- Service Processor menu
- Operating system service menu
- Service Focal Point on the HMC or vHMC with PowerVM

In the light path LED implementation, the system can clearly identify components for replacement by using specific component-level LEDs, and can also guide the servicer directly to the component by signaling (turning on solid) the amber system fault LED, enclosure fault LED, and component FRU fault LED. The servicer can also use the identify function to blink the FRU-level LED. When this function is activated, a roll-up to the blue enclosure locate and system locate LEDs will occur. These enclosure LEDs will turn on solid and can be used to follow the light path from the system to the enclosure and down to the specific FRU in the PowerVM environment.

First Failure Data Capture and error data analysis

First Failure Data Capture (FFDC) is a technique that helps ensure that when a fault is detected in a system, the root cause of the fault will be captured without the need to re-create the problem or run any sort of extending tracing or diagnostics program. For the vast majority of faults, a good FFDC design means that the root cause can also be detected automatically without servicer intervention.

FFDC information, error data analysis, and fault isolation are necessary to implement the advanced serviceability techniques that enable efficient service of the systems and to help determine the failing items.

In the rare absence of FFDC and Error Data Analysis, diagnostics are required to re-create the failure and determine the failing items.

Diagnostics

General diagnostic objectives are to detect and identify problems so they can be resolved quickly. Elements of IBM's diagnostics strategy include:

- Provide a common error code format equivalent to a system reference code with PowerVM, system reference number, checkpoint, or firmware error code.
- Provide fault detection and problem isolation procedures. Support remote connection ability to be used by the IBM Remote Support Center or IBM Designated Service.
- Provide interactive intelligence within the diagnostics with detailed online failure information while connected to IBM's back-end system.

Automatic diagnostics

The processor and memory FFDC technology is designed to perform without the need for recreate diagnostics nor require user intervention. Solid and intermittent errors are designed to be correctly detected and isolated at the time the failure occurs. Runtime and boot-time diagnostics fall into this category.

Stand-alone diagnostics with PowerVM

As the name implies, stand-alone or user-initiated diagnostics requires user intervention. The user must perform manual steps, including:

- Booting from the diagnostics CD, DVD, USB, or network
- Interactively selecting steps from a list of choices

Concurrent maintenance

The determination of whether a firmware release can be updated concurrently is identified in the readme information file that is released with the firmware. An HMC is required for the concurrent firmware update with PowerVM. In addition, concurrent maintenance of PCIe adapters is supported with PowerVM. Concurrent maintenance of the Operator Panel is supported through ASMI. Additional concurrent maintenance includes power supplies, fans, and HDD/SSD drives.

Service labels

Service providers use these labels to assist them in performing maintenance actions. Service labels are found in various formats and positions and are intended to transmit readily available information to the servicer during the repair process. Following are some of these service labels and their purpose:

- **Location diagrams:** Location diagrams are located on the system hardware, relating information regarding the placement of hardware components. Location diagrams may include location codes, drawings of physical locations, concurrent maintenance status, or other data pertinent to a repair. Location diagrams are especially useful when multiple components such as DIMMs, CPUs, processor books, fans, adapter cards, LEDs, and power supplies are installed.
- **Remove/replace procedures:** Service labels that contain remove/replace procedures are often found on a cover of the system or in other spots accessible to the servicer. These labels provide systematic procedures, including diagrams detailing how to remove or replace certain serviceable hardware components.
- **Arrows:** Numbered arrows are used to indicate the order of operation and the serviceability direction of components. Some serviceable parts such as latches, levers, and touch points need to be pulled or pushed in a certain direction and in a certain order for the mechanical mechanisms to engage or disengage. Arrows generally improve the ease of serviceability.

QR labels

QR labels are placed on the system to provide access to key service functions through a mobile device. When the QR label is scanned, it will go to a landing page specific to that server which contains many of the service functions of interest while physically located at the server. These include things such as installation and repair instructions, service diagrams, reference code look up, and so on.

Packaging for service

The following service enhancements are included in the physical packaging of the systems to facilitate service:

- **Color coding (touch points):** Blue-colored touch points delineate touchpoints on service components where the component can be safely handled for service actions such as removal or installation.
- **Tool-less design:** Selected IBM systems support tool-less or simple tool designs. These designs require no tools or simple tools such as flathead screw drivers to service the hardware components.
- **Positive retention:** Positive retention mechanisms help to assure proper connections between hardware components such as cables to connectors, and between two cards that attach to each other. Without positive retention, hardware components run the risk of becoming loose during shipping or installation, preventing a good electrical connection. Positive retention mechanisms like latches, levers, thumb-screws, pop Nylatches (U-clips), and cables are included to help prevent loose connections and aid in installing (seating) parts correctly. These positive retention items do not require tools.

Error handling and reporting

In the event of system hardware or environmentally induced failure, the system runtime error capture capability systematically analyzes the hardware error signature to determine the cause of failure. The analysis result will be stored in system NVRAM. When the system can be successfully restarted either manually or automatically, or if the system continues to operate, the error will be reported to the operating system. Hardware and software failures are recorded in the system log. When an HMC is attached in the PowerVM environment, an 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, forwards them to the SFP application, and notifies the system administrator.

The system has the ability to call home through the operating system to report platform-recoverable errors and errors associated with PCI adapters/devices.

In the HMC-managed environment, a call home service request will be initiated from the HMC and the pertinent failure data with service parts information and part locations will be 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 error log data related to the failure, are sent to IBM Service.

Live Partition Mobility

With PowerVM Live Partition Mobility (LPM), users can migrate an AIX, IBM i or Linux VM running on one POWER system to another POWER system without disrupting services. The migration transfers the entire system environment, including processor state, memory, attached virtual devices, and connected users. It provides continuous operating system and application availability during planned outages for repair of hardware and firmware faults. The Power System E950 server (and other servers using POWER9 technology processors with firmware level FW920 or above) supports secure LPM, whereby the VM image is encrypted and compressed prior to transfer. Secure LPM uses on-chip encryption and compression capabilities of the POWER9 processor for optimal performance.

Service processor

Diagnostic monitoring of recoverable error from the processor chipset is performed on the system processor itself, while the fatal diagnostic monitoring of the processor chipset is performed by the service processor.

Under PowerVM, the service processor supports surveillance of the connection to the HMC and to the system firmware (hypervisor). It also provides several remote power control options, environmental monitoring, reset, restart, 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

Call home refers to an automatic or manual call from a client location to the IBM support structure with error log data, server status, or other service-related information. Call home invokes the service organization in order for the appropriate service action to begin. Call home can be done through HMC or most non-HMC-managed systems through Electronic Service Agent running on top of the operating system. While configuring call home is optional, clients are encouraged to implement this feature in order to obtain service enhancements such as reduced problem determination and faster and potentially more accurate transmittal of error information. In general, using the call home feature can result in increased system availability. The Electronic Service Agent application can be configured for automated call home. See the next section for specific details on this application.

IBM Electronic Services

Electronic Service Agent and the IBM Electronic Services web portal comprise the IBM Electronic Services solution, which is dedicated to providing fast, exceptional support to IBM clients. IBM Electronic Service Agent is a no-charge tool that proactively monitors and reports hardware events such as system errors, performance issues, and inventory. Electronic Service Agent can help focus on the client's company business initiatives, save time, and spend less effort managing day-to-day IT maintenance issues.

System configuration and inventory information collected by Electronic Service Agent also can be viewed on the secure Electronic Services web portal and used to improve problem determination and resolution between the client and the IBM support team. As part of an increased focus to provide even better service to IBM

clients, Electronic Service Agent tool configuration and activation comes standard with the system. In support of this effort, a new HMC External Connectivity security whitepaper has been published, which describes data exchanges between the HMC and the IBM Service Delivery Center (SDC) and the methods and protocols for this exchange. To read the whitepaper and prepare for Electronic Service Agent installation, see the "Security" section at the [IBM Electronic Service Agent](#) website.

Select your country. Click "IBM Electronic Service Agent Connectivity Guide."

Benefits: increased uptime

Electronic Service Agent is designed to enhance the warranty and maintenance service by potentially providing faster hardware error reporting and uploading system information to IBM Support. This can optimize the time monitoring the symptoms, diagnosing the error, and manually calling IBM Support to open a problem record. And 24x7 monitoring and reporting means no more dependency on human intervention or off-hours client personnel when errors are encountered in the middle of the night.

Security: The Electronic Service Agent tool is designed to help secure the monitoring, reporting, and storing of the data at IBM. The Electronic Service Agent tool is designed to help securely transmit either through the internet (HTTPS or VPN) or modem to provide clients a single point of exit from their site. Communication is one way. Activating Electronic Service Agent does not enable IBM to call into a client's system.

For additional information, see the [IBM Electronic Service Agent](#) website.

More accurate reporting

Because system information and error logs are automatically uploaded to the IBM Support Center in conjunction with the service request, clients are not required 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

By using the IBMid entered during activation, clients can view system and support information in the "My Systems" and "Premium Search" sections of the Electronic Services website.

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 Premium Search functions make it even easier for Electronic Service Agent-enabled clients to track system inventory and find pertinent fixes.

My Systems provides valuable reports of installed hardware and software using information collected from the systems by IBM Electronic Service Agent. Reports are available for any system associated with the client's IBMid. Premium Search combines the function of search and the value of Electronic Service Agent information, providing advanced search of the technical support knowledgebase. Using Premium Search and the Service Agent information that has been collected from the system, clients are able to see search results that apply specifically to their systems.

For more information on how to utilize the power of IBM Electronic Services, see the following website or contact an [IBM Systems Services Representative](#).

IBM Proactive Support for Power Systems

Clients have found significant value in IBM's Proactive Support offerings on mission-critical systems, as this provides personalized support, proactive recommendations,

and accelerated response times versus standard support. As a result, IBM is including the IBM Proactive Support in the default configuration for all Mid-Range and Enterprise IBM POWER9 Systems - for IBM AIX, IBM i, Linux, and SAP HANA workloads. Other configurations are also available.

Statement of general direction

IBM Power System E950 (9040-MR9) for SAP HANA Production Use

IBM intends to support SAP HANA on the IBM Power System E950 (9040-MR9) in production mode, with the following Linux operating systems, following near-term certification of the environment.

- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 7 for Power LE version 7.5, or later
- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 11 Service Pack 4, or later
- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 12 Service Pack 3, or later

Statements by IBM regarding its plans, directions, and intent are subject to change or withdrawal without notice at the sole discretion of IBM. Information regarding potential future products is intended to outline general product direction and should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for IBM products remain at the sole discretion of IBM.

Product number

The following are newly announced features on the specific models of the IBM Power Systems 9040 machine type:

Machine Description	Feature	type	Model	number
IBM Power System E950		9040	MR9	
EMEA Bulk MES Indicator		9040	MR9	0004
One CSC Billing Unit		9040	MR9	0010
Ten CSC Billing Units		9040	MR9	0011
Special Manufacturing Operations Indicator		9040	MR9	0098
AIX Partition Specify		9040	MR9	0265
Linux Partition Specify		9040	MR9	0266
V.24/EIA232 6.1m (20-Ft) PCI Cable		9040	MR9	0348
V.35 6.1m (20-Ft) PCI Cable		9040	MR9	0353
X.21 6.1m (20-Ft) PCI Cable		9040	MR9	0359
V.24/EIA232 20-Ft. PCI Cable with M3		9040	MR9	0368
UPS Factory Integration Specify		9040	MR9	0373
HMC Factory Integration Specify		9040	MR9	0374
Display Factory Integration Specify		9040	MR9	0375
Reserve Rack Space for UPS		9040	MR9	0376
Reserve Rack Space for HMC		9040	MR9	0377
Reserve Rack Space for Display		9040	MR9	0378
Customer Specified Placement		9040	MR9	0456
19 inch, 1.8 meter high rack		9040	MR9	0551
19 inch, 2.0 meter high rack		9040	MR9	0553
Rack Filler Panel Kit		9040	MR9	0599
Load Source Not in CEC		9040	MR9	0719
SAN Load Source Specify		9040	MR9	0837
Modem Cable - Austria		9040	MR9	1010
Modem Cable - Belgium		9040	MR9	1011
Modem Cable - Africa		9040	MR9	1012
Modem Cable - Italy		9040	MR9	1014
Modem Cable - France		9040	MR9	1015
Modem Cable - Germany		9040	MR9	1016

Modem Cable - UK	9040	MR9	1017
Modem Cable - Iceland/Sweden	9040	MR9	1018
Modem Cable - Fin/Nor	9040	MR9	1021
Modem Cable - Netherlands	9040	MR9	1022
Modem Cable - Swiss	9040	MR9	1023
Modem Cable - Denmark	9040	MR9	1024
Modem Cable - US/Canada and General Use	9040	MR9	1025
USB 500 GB Removable Disk Drive	9040	MR9	1107
3m, Blue Cat5e Cable	9040	MR9	1111
10m, Blue Cat5e Cable	9040	MR9	1112
25m, Blue Cat5e Cable	9040	MR9	1113
Custom Service Specify, Rochester Minn, USA	9040	MR9	1140
Quantity 150 of #1964	9040	MR9	1818
Quantity 150 of #1953	9040	MR9	1929
300GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)	9040	MR9	1953
600GB 10k RPM SAS SFF-2 Disk Drive (AIX/Linux)	9040	MR9	1964
Primary OS - AIX	9040	MR9	2146
Primary OS - Linux	9040	MR9	2147
2M LC-SC 50 Micron Fiber Converter Cable	9040	MR9	2456
2M LC-SC 62.5 Micron Fiber Converter Cable	9040	MR9	2459
3M Asynchronous Terminal/Printer Cable EIA-232	9040	MR9	2934
Asynchronous Cable EIA-232/V.24 3M	9040	MR9	2936
Serial-to-Serial Port Cable for Drawer/Drawer- 3.7M	9040	MR9	3124
Serial-to-Serial Port Cable for Rack/Rack- 8M	9040	MR9	3125
Widescreen LCD Monitor	9040	MR9	3632
0.3M Serial Port Converter Cable, 9-Pin to 25-Pin	9040	MR9	3925
Serial Port Null Modem Cable, 9-pin to 9-pin, 3.7M	9040	MR9	3927
Serial Port Null Modem Cable, 9-pin to 9-pin, 10M	9040	MR9	3928
System Serial Port Converter Cable	9040	MR9	3930
1.8 M (6-ft) Extender Cable for Displays (15-pin D-shell to 15-pin D-shell)	9040	MR9	4242
Extender Cable - USB Keyboards, 1.8M	9040	MR9	4256
VGA to DVI Connection Converter	9040	MR9	4276
Power Cord To PDU/UPS, (100-240V/16A)	9040	MR9	4558
Rack Integration Services: BP only	9040	MR9	4648
Rack Integration Services	9040	MR9	4649
Rack Indicator- Not Factory Integrated	9040	MR9	4650
Software Preload Required	9040	MR9	5000
PCIe2 8Gb 4-port Fibre Channel Adapter	9040	MR9	5729
8 Gigabit PCI Express Dual Port Fibre Channel Adapter	9040	MR9	5735
PCIe2 4-port 1GbE Adapter	9040	MR9	5899
Opt Front Door for 1.8m Rack	9040	MR9	6068
Opt Front Door for 2.0m Rack	9040	MR9	6069
1.8m Rack Acoustic Doors	9040	MR9	6248
2.0m Rack Acoustic Doors	9040	MR9	6249
1.8m Rack Trim Kit	9040	MR9	6263
2.0m Rack Trim Kit	9040	MR9	6272
Power Cord 4.3m (14-ft), Drawer to IBM PDU (250V/ 10A)	9040	MR9	6458
Power Cord 4.3m (14-ft), Drawer To OEM PDU (125V, 15A)	9040	MR9	6460
Power Cord 4.3m (14-ft), Drawer to wall/OEM PDU (250V/15A) U. S.	9040	MR9	6469
Power Cord 1.8m (6-ft), Drawer to wall (125V/15A)	9040	MR9	6470
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/10A)	9040	MR9	6471
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/16A)	9040	MR9	6472
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/10A)	9040	MR9	6473
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/13A)	9040	MR9	6474
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/16A)	9040	MR9	6475
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9040	MR9	6476
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/16A)	9040	MR9	6477
Power Cord 2.7 M(9-foot), To wall/OEM PDU, (250V, 16A)	9040	MR9	6478
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU,			

(125V/15A or 250V/10A)	9040	MR9	6488
4.3m (14-Ft) 3PH/24A 380-415V Power Cord	9040	MR9	6489
4.3m (14-Ft) 1PH/63A 200-240V Power Cord	9040	MR9	6491
4.3m (14-Ft) 1PH/48-60A 200-240V Power Cord	9040	MR9	6492
Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)	9040	MR9	6493
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9040	MR9	6494
Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 10A)	9040	MR9	6496
Power Cable - Drawer to IBM PDU, 200-240v/10A	9040	MR9	6577
Optional Rack Security Kit	9040	MR9	6580
Power Cord 2.7M (9-foot), To Wall/OEM PDU, (125V, 15A)	9040	MR9	6651
4.3m (14-Ft) 3PH/16A 380-415V Power Cord	9040	MR9	6653
4.3m (14-Ft) 1PH/24-30A Pwr Cord	9040	MR9	6654
4.3m (14-Ft) 1PH/24-30A WR Pwr Cord	9040	MR9	6655
4.3m (14-Ft) 1PH/24A Power Cord	9040	MR9	6656
4.3m (14-Ft) 1PH/32A Power Cord	9040	MR9	6657
4.3m (14-Ft) 1PH/24A Pwr Cd-Korea	9040	MR9	6658
Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 15A)	9040	MR9	6659
Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (125V/15A)	9040	MR9	6660
Power Cord 2.8m (9.2-ft), Drawer to IBM PDU, (250V/10A)	9040	MR9	6665
4.3m (14-Ft) 3PH/32A 380-415V Power Cord-Australia	9040	MR9	6667
Power Cord 4.3M (14-foot), Drawer to OEM PDU, (250V, 15A)	9040	MR9	6669
Power Cord 2.7M (9-foot), Drawer to IBM PDU, 250V/10A	9040	MR9	6671
Power Cord 2M (6.5-foot), Drawer to IBM PDU, 250V/10A	9040	MR9	6672
Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)	9040	MR9	6680
Environmental Monitoring Probe	9040	MR9	7118
OEM (Generic) Indicator	9040	MR9	7770
OEM (GROUPE BULL) Indicator	9040	MR9	7773
OEM (Hitachi) Indicator	9040	MR9	7775
OEM Publications for IBM Logo Product	9040	MR9	7779
Indicator for Inspur Power servers	9040	MR9	777A
Ethernet Cable, 15m, Hardware Management Console to System Unit	9040	MR9	7802
Linux Software Preinstall	9040	MR9	8143
Linux Software Preinstall (Business Partners)	9040	MR9	8144
USB Mouse	9040	MR9	8845
Order Routing Indicator- System Plant	9040	MR9	9169
Language Group Specify - US English	9040	MR9	9300
New AIX License Core Counter	9040	MR9	9440
New Red Hat License Core Counter	9040	MR9	9442
New SUSE License Core Counter	9040	MR9	9443
Other AIX License Core Counter	9040	MR9	9444
Other Linux License Core Counter	9040	MR9	9445
3rd Party Linux License Core Counter	9040	MR9	9446
VIOS Core Counter	9040	MR9	9447
Other License Core Counter	9040	MR9	9449
Ubuntu Linux License Core Counter	9040	MR9	9450
Month Indicator	9040	MR9	9461
Day Indicator	9040	MR9	9462
Hour Indicator	9040	MR9	9463
Minute Indicator	9040	MR9	9464
Qty Indicator	9040	MR9	9465
Countable Member Indicator	9040	MR9	9466
Language Group Specify - Dutch	9040	MR9	9700
Language Group Specify - French	9040	MR9	9703
Language Group Specify - German	9040	MR9	9704
Language Group Specify - Polish	9040	MR9	9705
Language Group Specify - Norwegian	9040	MR9	9706
Language Group Specify - Portuguese	9040	MR9	9707
Language Group Specify - Spanish	9040	MR9	9708
Language Group Specify - Italian	9040	MR9	9711
Language Group Specify - Canadian French	9040	MR9	9712
Language Group Specify - Japanese	9040	MR9	9714

Language Group Specify - Traditional Chinese (Taiwan)	9040	MR9	9715
Language Group Specify - Korean	9040	MR9	9716
Language Group Specify - Turkish	9040	MR9	9718
Language Group Specify - Hungarian	9040	MR9	9719
Language Group Specify - Slovakian	9040	MR9	9720
Language Group Specify - Russian	9040	MR9	9721
Language Group Specify - Simplified Chinese (PRC)	9040	MR9	9722
Language Group Specify - Czech	9040	MR9	9724
Language Group Specify - Romanian	9040	MR9	9725
Language Group Specify - Croatian	9040	MR9	9726
Language Group Specify - Slovenian	9040	MR9	9727
Language Group Specify - Brazilian Portuguese	9040	MR9	9728
Language Group Specify - Thai	9040	MR9	9729
QSFP+ 40GbE Base-SR Transceiver	9040	MR9	EB27
1m (3.3-ft), IBM Passive QSFP+ to QSFP+ Cable (DAC)	9040	MR9	EB2B
3m (9.8-ft), IBM Passive QSFP+ to QSFP+ Cable (DAC)	9040	MR9	EB2H
10m (30.3-ft), IBM Passive QSFP+ MTP Optical Cable	9040	MR9	EB2J
30m (90.3-ft), IBM Passive QSFP+ MTP Optical Cable	9040	MR9	EB2K
Lift Tool	9040	MR9	EB2Z
Power Supply - 2000W for Server (200-240 VAC)	9040	MR9	EB3M
Lift tool based on GenieLift GL-8 (standard)	9040	MR9	EB3Z
10GbE Optical Transceiver SFP+ SR	9040	MR9	EB46
25GbE Optical Transceiver SFP28	9040	MR9	EB47
0.5m SFP28/25GbE copper Cable	9040	MR9	EB4J
1.0m SFP28/25GbE copper Cable	9040	MR9	EB4K
1.5m SFP28/25GbE copper Cable	9040	MR9	EB4L
2.0m SFP28/25GbE copper Cable	9040	MR9	EB4M
2.0m QSFP28/100GbE copper split Cable to SFP28 4x25GbE	9040	MR9	EB4P
Service wedge shelf tool kit for EB3Z	9040	MR9	EB4Z
0.5m EDR IB Copper Cable QSFP28	9040	MR9	EB50
1.0m EDR IB Copper Cable QSFP28	9040	MR9	EB51
2.0M EDR IB Copper Cable QSFP28	9040	MR9	EB52
1.5M EDR IB Copper Cable QSFP28	9040	MR9	EB54
100GbE Optical Transceiver QSFP28	9040	MR9	EB59
3M EDR IB Optical Cable QSFP28	9040	MR9	EB5A
5M EDR IB Optical Cable QSFP28	9040	MR9	EB5B
10M EDR IB Optical Cable QSFP28	9040	MR9	EB5C
15M EDR IB Optical Cable QSFP28	9040	MR9	EB5D
20M EDR IB optical Cable QSFP28	9040	MR9	EB5E
30M EDR IB optical Cable QSFP28	9040	MR9	EB5F
50M EDR IB Optical Cable QSFP28	9040	MR9	EB5G
100M EDR IB optical cable QSFP28	9040	MR9	EB5H
0.5M 100GbE Copper Cable QSFP28	9040	MR9	EB5J
1.0M 100GbE Copper Cable QSFP28	9040	MR9	EB5K
1.5M 100GbE Copper Cable QSFP28	9040	MR9	EB5L
2.0M 100GbE Copper Cable QSFP28	9040	MR9	EB5M
3M 100GbE Optical cable QSFP28 (AOC)	9040	MR9	EB5R
5M 100GbE Optical Cable QSFP28 (AOC)	9040	MR9	EB5S
10M 100GbE optical Cable QSFP28 (AOC)	9040	MR9	EB5T
15M 100GbE optical Cable QSFP28 (AOC)	9040	MR9	EB5U
20M 100GbE Optical Cable QSFP28 (AOC)	9040	MR9	EB5V
30M 100GbE optical Cable QSFP28 (AOC)	9040	MR9	EB5W
50M 100GbE optical Cable QSFP28 (AOC)	9040	MR9	EB5X
100M 100GbE optical Cable QSFP28 (AOC)	9040	MR9	EB5Y
Rack Front Door (Black)	9040	MR9	EC01
Rack Rear Door	9040	MR9	EC02
Rack Side Cover	9040	MR9	EC03
Rack Suite Attachment Kit	9040	MR9	EC04
Slim Rear Acoustic Door	9040	MR9	EC07
Slim Front Acoustic Door	9040	MR9	EC08
Rear Door Heat Exchanger for 2.0 Meter Slim Rack	9040	MR9	EC15
PCIe3 2-Port 10Gb NIC&ROCE SR/Cu Adapter	9040	MR9	EC2S
PCIe3 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter	9040	MR9	EC2U
Mainstream 800 GB SSD NVMe U.2 module	9040	MR9	EC5J
Mainstream 1.6 TB SSD NVMe U.2 module	9040	MR9	EC5K
Mainstream 3.2 TB SSD NVMe U.2 module	9040	MR9	EC5L
PCIe4 2-port 100Gb ROCE EN adapter	9040	MR9	EC66
0.6m (2.0-ft), Blue CAT5 Ethernet Cable	9040	MR9	ECB0

1.5m (4.9-ft), Blue CAT5 Ethernet Cable	9040	MR9	ECB2
SAS X Cable 3m - HD Narrow 6Gb 2-Adapters to Enclosure	9040	MR9	ECBJ
SAS X Cable 6m - HD Narrow 6Gb 2-Adapters to Enclosure	9040	MR9	ECBK
SAS X Cable 10m - HD Narrow 6Gb 2-Adapters to Enclosure	9040	MR9	ECBL
SAS X Cable 15m - HD Narrow 3Gb 2-Adapters to Enclosure	9040	MR9	ECBM
SAS YO Cable 1.5m - HD Narrow 6Gb Adapter to Enclosure	9040	MR9	ECBT
SAS YO Cable 3m - HD Narrow 6Gb Adapter to Enclosure	9040	MR9	ECBU
SAS YO Cable 6m - HD Narrow 6Gb Adapter to Enclosure	9040	MR9	ECBV
SAS YO Cable 10m - HD Narrow 6Gb Adapter to Enclosure	9040	MR9	ECBW
SAS YO Cable 15m - HD Narrow 3Gb Adapter to Enclosure	9040	MR9	ECBX
SAS AE1 Cable 4m - HD Narrow 6Gb Adapter to Enclosure	9040	MR9	ECBY
SAS YE1 Cable 3m - HD Narrow 6Gb Adapter to Enclosure	9040	MR9	ECBZ
SAS AA Cable 0.6m - HD Narrow 6Gb Adapter to Adapter	9040	MR9	ECC0
SAS AA Cable 1.5m - HD Narrow 6Gb Adapter to Adapter	9040	MR9	ECC2
SAS AA Cable 3m - HD Narrow 6Gb Adapter to Adapter	9040	MR9	ECC3
SAS AA Cable 6m - HD Narrow 6Gb Adapter to Adapter	9040	MR9	ECC4
3M Optical Cable Pair for PCIe3 Expansion Drawer	9040	MR9	ECC7
10M Optical Cable Pair for PCIe3 Expansion Drawer	9040	MR9	ECC8
System Port Converter Cable for UPS	9040	MR9	ECCF
3M Copper CXP Cable Pair for PCIe3 Expansion Drawer	9040	MR9	ECCS
3.0M SAS X12 Cable (Two Adapter to Enclosure)	9040	MR9	ECDJ
4.5M SAS X12 Active Optical Cable (Two Adapter to Enclosure)	9040	MR9	ECDK
10M SAS X12 Active Optical Cable (Two Adapter to Enclosure)	9040	MR9	ECDL
1.5M SAS Y012 Cable (Adapter to Enclosure)	9040	MR9	ECDT
3.0M SAS Y012 Cable (Adapter to Enclosure)	9040	MR9	ECDU
4.5M SAS Y012 Active Optical Cable (Adapter to Enclosure)	9040	MR9	ECDV
10M SAS Y012 Active Optical Cable (Adapter to Enclosure)	9040	MR9	ECDW
0.6M SAS AA12 Cable (Adapter to Adapter)	9040	MR9	ECE0
3.0M SAS AA12 Cable	9040	MR9	ECE3
4.5M SAS AA12 Active Optical Cable (Adapter to Adapter)	9040	MR9	ECE4
Cloud Private Solution	9040	MR9	ECPO
2.0 Meter Slim Rack	9040	MR9	ECR0
Rack Front Door High-End appearance	9040	MR9	ECRF
Rack Rear Door Black	9040	MR9	ECRG
Rack Side Cover	9040	MR9	ECRJ
Rack Rear Extension 5-In	9040	MR9	ECRK
Rack Front Door for Rack (Black/Flat)	9040	MR9	ECRM
Custom Service Specify, Montpellier, France	9040	MR9	ECSF
NeuCloud Indicator/Specify	9040	MR9	ECSJ
Custom Service Specify, Mexico	9040	MR9	ECSM
Custom Service Specify, Poughkeepsie, USA	9040	MR9	ECSP
Optical wrap Plug	9040	MR9	ECW0
SAP HANA TRACKING FEATURE	9040	MR9	EHKV
Boot Drive / Load Source in EXP12SX Specify (in #ESLL or #ELLL)	9040	MR9	EHR1
Boot Drive / Load Source in EXP24SX Specify (in #ESLS or #ELLS)	9040	MR9	EHR2
SSD Placement Indicator - #ESLS/#ELLS	9040	MR9	EHS2
PCIe3 Optical Cable Adapter for PCIe3 Expansion Drawer	9040	MR9	EJ08
Storage Backplane with Zero DASD 8 SAS 2.5" HDD/SDD Controllers	9040	MR9	EJ0B
PCIe3 RAID SAS Adapter Quad-port 6Gb x8	9040	MR9	EJ0J

PCIe3 RAID SAS Adapter Quad-port 6Gb x8 for MR9	9040	MR9	EJ0K
PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8	9040	MR9	EJ10
PCIe3 12GB Cache RAID PLUS SAS Adapter Quad-port 6Gb x8	9040	MR9	EJ14
PCIe3 Crypto Coprocessor BSC-Gen3 4767	9040	MR9	EJ33
Storage Backplane Base DASD 8 SAS 2.5" HDD/SDD Controllers	9040	MR9	EJBB
Non-paired Indicator EJ14 PCIe SAS RAID+ Adapter Storage Backplane Split DASD 8 SAS 2.5" HDD/SDD Controllers	9040	MR9	EJRL
Specify Mode-1 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP12SX #ESLL/ELLL	9040	MR9	EJSB
Specify Mode-1 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP12SX #ESLL/ELLL	9040	MR9	EJV1
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP12SX #ESLL/ELLL	9040	MR9	EJV2
Specify Mode-2 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP12SX #ESLL/ELLL	9040	MR9	EJV3
Specify Mode-4 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP12SX #ESLL/ELLL	9040	MR9	EJV4
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP12SX #ESLL/ELLL	9040	MR9	EJV5
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP12SX #ESLL/ELLL	9040	MR9	EJV6
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP12SX #ESLL/ELLL	9040	MR9	EJV7
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP12SX #ESLL/ELLL	9040	MR9	EJVA
Specify Mode-4 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP12SX #ESLL/ELLL	9040	MR9	EJVB
Specify Mode-4 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP12SX #ESLL/ELLL	9040	MR9	EJVC
Specify Mode-4 & (3)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP12SX #ESLL/ELLL	9040	MR9	EJVD
Specify Mode-1 & (2)EJ14 & (2)Y012 for EXP12SX #ESLL/ELLL	9040	MR9	EJVE
Specify Mode-2 & (1)EJ0K & (2)Y012 for EXP12SX #ESLL/ELLL	9040	MR9	EJVF
Specify Mode-1 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP24SX #ESLS/ELLS	9040	MR9	EJVV
Specify Mode-1 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9040	MR9	EJW1
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJW2
Specify Mode-2 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJW3
Specify Mode-4 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJW4
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9040	MR9	EJW5
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9040	MR9	EJW6
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP24SX #ESLS/ELLS	9040	MR9	EJW7
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJWA
Specify Mode-4 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJWB
Specify Mode-4 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJWC
Specify Mode-4 & (3)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJWD
Specify Mode-1 & (2)EJ14 & (2)Y012 for EXP24SX #ESLS/ELLS	9040	MR9	EJWE
Specify Mode-2 & (2)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJWF
Specify Mode-2 & (2)EJ14 & (1)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJWG
Specify Mode-2 & (4)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS	9040	MR9	EJWH
Specify Mode-1 & (2)EJ0K & (2)Y012G for EXP24SX #ESLS/ELLS	9040	MR9	EJWJ
1-core Linux Processor Activation for #EPWR	9040	MR9	EJWW
1-core Linux Processor Activation for #EPWT	9040	MR9	ELBG
			ELBH

PowerVM for Linux indicator	9040	MR9	ELBJ
1-core Linux Processor Activation for #EPWS	9040	MR9	ELBP
1-core Linux Processor Activation for #EPWY	9040	MR9	ELBR
PDU Access Cord 0.38m	9040	MR9	ELC0
Power Cable - Drawer to IBM PDU (250V/10A)	9040	MR9	ELC5
Power IFL Memory Activation	9040	MR9	ELNP
Memory Riser Card	9040	MR9	EM03
8 GB DDR4 Memory	9040	MR9	EM6A
16 GB DDR4 Memory	9040	MR9	EM6B
32 GB DDR4 Memory	9040	MR9	EM6C
64 GB DDR4 Memory	9040	MR9	EM6D
128 GB DDR4 Memory	9040	MR9	EM6E
Active Memory Mirroring	9040	MR9	EM81
90 Days Elastic CoD Memory Enablement	9040	MR9	EM9U
Power Active Memory Expansion	9040	MR9	EMAM
1GB Memory Activation	9040	MR9	EMAP
Quantity of 100 1GB Memory Activations	9040	MR9	EMAQ
512 GB Linux Memory Activations for MR9	9040	MR9	EMBE
VRM DDR4 Memory for MR9	9040	MR9	EMEF
8 GB-Day Billing for Elastic CoD Memory	9040	MR9	EMJE
800 GB-Day Billing for Elastic CoD Memory	9040	MR9	EMJF
PCIe Gen3 I/O Expansion Drawer	9040	MR9	EMX0
AC Power Supply Conduit for PCIe3 Expansion Drawer	9040	MR9	EMXA
PCIe3 6-slot Fanout Module for PCIe3 Expansion Drawer	9040	MR9	EMXF
PCIe3 6-slot Fanout Module for PCIe3 Expansion Drawer	9040	MR9	EMXG
1m (3.3-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper	9040	MR9	EN01
3m (9.8-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper	9040	MR9	EN02
5m (16.4-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper	9040	MR9	EN03
PCIe3 16Gb 2-port Fibre Channel Adapter	9040	MR9	EN0A
PCIe3 4-port (10Gb FCoE & 1GbE) SR&RJ45	9040	MR9	EN0H
PCIe3 4-port (10Gb FCoE & 1GbE) SFP+Copper&RJ45	9040	MR9	EN0K
PCIe2 4-Port (10Gb+1GbE) SR+RJ45 Adapter	9040	MR9	EN0S
PCIe2 4-port (10Gb+1GbE) Copper SFP+RJ45 Adapter	9040	MR9	EN0U
PCIe2 2-port 10/1GbE BaseT RJ45 Adapter	9040	MR9	EN0W
PCIe2 8Gb 4-port Fibre Channel Adapter	9040	MR9	EN12
PCIe3 4-port 10GbE SR Adapter	9040	MR9	EN15
PCIe3 32Gb 2-port Fibre Channel Adapter	9040	MR9	EN1A
PCIe3 16Gb 4-port Fibre Channel Adapter	9040	MR9	EN1C
90 Days Elastic CoD Processor Core Enablement	9040	MR9	EP9T
1 Proc-day Elastic billing for #EPWR/EPWK	9040	MR9	EPN0
100 Proc-day Elastic COD billing #EPWR/EPWK AIX	9040	MR9	EPN1
100 Proc-mins Utility COD billing #EPWR/EPWK	9040	MR9	EPN2
1-core Processor Activation for #EPWY	9040	MR9	EPN3
1 Proc-day Elastic billing for #EPWS/EPWL	9040	MR9	EPN5
100 Proc-day Elastic COD billing #EPWS/EPWL AIX	9040	MR9	EPN6
100 Proc-mins Utility COD billing #EPWS/EPWL	9040	MR9	EPN7
1 Proc-day Elastic billing for #EPWY/EPWZ	9040	MR9	EPN8
100 Proc-day Elastic COD billing #EPWY/EPWZ AIX	9040	MR9	EPN9
1 Proc-day Elastic billing for #EPWT/EPWM	9040	MR9	EPNK
100 Proc-day Elastic COD billing #EPWT/EPWM AIX	9040	MR9	EPNL
100 Proc-mins Utility COD billing #EPWT/EPWM	9040	MR9	EPNM
100 Proc-mins Utility COD billing #EPWY/EPWZ	9040	MR9	EPNN
Horizontal PDU Mounting Hardware	9040	MR9	EPTH
High Function 9xC19 PDU: Switched, Monitoring	9040	MR9	EPTJ
High Function 9xC19 PDU 3-Phase: Switched, Monitoring	9040	MR9	EPTL
High Function 12xC13 PDU: Switched, Monitoring	9040	MR9	EPTN
High Function 12xC13 PDU 3-Phase: Switched, Monitoring	9040	MR9	EPTQ
PowerVM - Enterprise Edition	9040	MR9	EPVV
8-core Typical 3.6 to 3.8 GHZ (max) processor	9040	MR9	EPWR
10-core Typical 3.4 to 3.8 GHZ (max) processor	9040	MR9	EPWS
12-core Typical 3.15 to 3.8 GHZ (max) processor	9040	MR9	EPWT
1-core Processor Activation for #EPWR	9040	MR9	EPWV
1-core Processor Activation for #EPWS	9040	MR9	EPWW
1W Processor activation for #EPWT	9040	MR9	EPWX
11-core Typical 3.2 to 3.8 GHZ (max) processor	9040	MR9	EPWY
Quantity 150 of #ES62 3.86-4.0 TB 7200 rpm 4k			

LFF-1 Disk	9040	MR9	EQ62
Quantity 150 of #ES64 7.72-8.0 TB 7200 rpm 4k			
LFF-1 Disk	9040	MR9	EQ64
Qty 150 of #6577	9040	MR9	EQ77
Quantity 150 of #ES78 387GB SFF-2 SSD 5xx	9040	MR9	EQ78
Quantity 150 of #ES7E 775GB SFF-2 SSD 5xx	9040	MR9	EQ7E
Quantity 150 of #ES80 1.9TB SFF-2 SSD 4k	9040	MR9	EQ80
Quantity 150 of #ES85 387GB SFF-2 SSD 4k	9040	MR9	EQ85
Quantity 150 of #ES8C 775GB SFF-2 SSD 4k	9040	MR9	EQ8C
Quantity 150 of #ES8F 1.55TB SFF-2 SSD 4k	9040	MR9	EQ8F
Quantity 150 of #ES8Y 931GB SFF-2 SSD 4k	9040	MR9	EQ8Y
Quantity 150 of ES96 1.86TB SFF-2 SSD 4k	9040	MR9	EQ96
Quantity 150 of #ESE7 3.72TB SFF-2 SSD 4k	9040	MR9	EQE7
Quantity 150 of #ESEV (600GB 10k SFF-2)	9040	MR9	EQEV
Quantity 150 of #ESEZ (300GB SFF-2)	9040	MR9	EQEZ
Quantity 150 of #ESF3 (1.2TB 10k SFF-2)	9040	MR9	EQF3
Quantity 150 of #ESFP (600GB SFF-2)	9040	MR9	EQFP
Quantity 150 of #ESFT (1.8TB 10k SFF-2)	9040	MR9	EQFT
Quantity 150 of #ESG5 (387GB SAS 5xx)	9040	MR9	EQG5
Quantity 150 of #ESGB (387GB SAS 4k)	9040	MR9	EQGB
Quantity 150 of #ESGF (775GB SAS 5xx)	9040	MR9	EQGF
Quantity 150 of #ESGK (775GB SAS 4k)	9040	MR9	EQGK
Quantity 150 of #ESGP (1.55TB SAS 4k)	9040	MR9	EQGP
42U Slim Rack	9040	MR9	ER05
RFID Tags for Servers, Compute Nodes, Chassis, Racks, and HMCs	9040	MR9	ERF1
Rear rack extension	9040	MR9	ERG0
Optional Origami Front Door for 2.0m Rack	9040	MR9	ERG7
Quantity 150 of #ESHJ 931 GB SSD 4k SFF-2	9040	MR9	ERHJ
Quantity 150 of #ESHL 1.86 TB SSD 4k SFF-2	9040	MR9	ERHL
Quantity 150 of #ESHN 7.45 TB SSD 4k SFF-2	9040	MR9	ERHN
Quantity 150 of #ESM8 3.72 TB SSD 4k SFF-2	9040	MR9	ERM8
3.86-4.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	9040	MR9	ES62
7.72-8.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux)	9040	MR9	ES64
387GB SFF-2 SSD 5xx eMLC4 for AIX/Linux	9040	MR9	ES78
775GB SFF-2 SSD 5xx eMLC4 for AIX/Linux	9040	MR9	ES7E
387GB SFF-3 SSD 5xx eMLC4 for AIX/Linux	9040	MR9	ES7K
775GB SFF-3 SSD 5xx eMLC4 for AIX/Linux	9040	MR9	ES7P
1.9TB Read Intensive SAS 4k SFF-2 SSD for AIX/ Linux	9040	MR9	ES80
931GB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ES83
387GB SFF-2 SSD 4k eMLC4 for AIX/Linux	9040	MR9	ES85
775GB SFF-2 SSD 4k eMLC4 for AIX/Linux	9040	MR9	ES8C
1.55TB SFF-2 SSD 4k eMLC4 for AIX/Linux	9040	MR9	ES8F
1.9TB Read Intensive SAS 4k SFF-3 SSD for AIX/ Linux	9040	MR9	ES8J
387GB SFF-3 SSD 4k eMLC4 for AIX/Linux	9040	MR9	ES8N
775GB SFF-3 SSD 4k eMLC4 for AIX/Linux	9040	MR9	ES8Q
1.55TB SFF-3 SSD 4k eMLC4 for AIX/Linux	9040	MR9	ES8V
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ES8Y
1.86TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ES92
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ES96
S&H - No Charge	9040	MR9	ESC0
S&H	9040	MR9	ESC7
600GB 10K RPM SAS SFF-3 Disk Drive (AIX/Linux)	9040	MR9	ESD5
300GB 15K RPM SAS SFF-3 Disk Drive (AIX/Linux)	9040	MR9	ESDB
3.72TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESE1
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESE7
600GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	9040	MR9	ESEV
300GB 15K RPM SAS SFF-2 4K Block - 4096 Disk Drive	9040	MR9	ESEZ
1.2TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	9040	MR9	ESF3
600GB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	9040	MR9	ESF5
1.2TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	9040	MR9	ESF9
300GB 15K RPM SAS SFF-3 4K Block - 4096 Disk Drive	9040	MR9	ESFB
600GB 15K RPM SAS SFF-3 4K Block - 4096 Disk Drive	9040	MR9	ESFF

600GB 15K RPM SAS SFF-2 4K Block - 4096 Disk Drive	9040	MR9	ESFP
1.8TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	9040	MR9	ESFT
1.8TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096	9040	MR9	ESFV
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9040	MR9	ESG5
387GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	9040	MR9	ESG9
387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESGB
387GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESGD
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9040	MR9	ESGF
775GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux	9040	MR9	ESGH
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESGK
775GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESGM
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESGP
1.55TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESGR
931 GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESHJ
1.86 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESHL
7.45 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESHN
931 GB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESHS
1.86 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESHU
7.45 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESHW
Specify AC Power Supply for EXP12SX/EXP24SX			
Storage Enclosure	9040	MR9	ESLA
EXP12SX SAS Storage Enclosure	9040	MR9	ESLL
EXP24SX SAS Storage Enclosure	9040	MR9	ESLS
3.72 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9040	MR9	ESM8
3.72 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux	9040	MR9	ESMQ
300GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	9040	MR9	ESNK
300GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	9040	MR9	ESNM
600GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)	9040	MR9	ESNP
600GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	9040	MR9	ESNR
Quantity 150 of #ESNM (300GB 15k SFF-2)	9040	MR9	ESPM
Quantity 150 of #ESNR (600GB 15k SFF-2)	9040	MR9	ESPR
1TB Removable Disk Drive Cartridge	9040	MR9	EU01
RDX USB External Docking Station for Removable Disk Cartridge	9040	MR9	EU04
RDX 320 GB Removable Disk Drive	9040	MR9	EU08
1.5TB Removable Disk Drive Cartridge	9040	MR9	EU15
Order Placed Indicator	9040	MR9	EU29
2TB Removable Disk Drive Cartridge (RDX)	9040	MR9	EU2T
RDX USB External Docking Station	9040	MR9	EUA4
Standalone USB DVD drive w/cable	9040	MR9	EUA5
Core Use HW Feature	9040	MR9	EUC6
Core Use HW Feature 10X	9040	MR9	EUC7
5000 Power to Cloud Reward points	9040	MR9	SVPC

The following are newly announced features on the specific models of the IBM Power Systems 7014 and 7965 machine type:

New features available August 17, 2018

Machine Description	Feature	type	Model	number
Rack Content Specify	9040-MR9 -4EIA	7014	T42	ERCZ
		7965	S42	

Publications

Power Systems hardware documentation provides you with the following topical information:

- Licenses, notices, safety, and warranty information
- Planning for the system
- Installing and configuring the system
- Troubleshooting, service, and support
- Installing, configuring, and managing consoles, terminals, and interfaces

- Installing operating systems
- Creating a virtual computing environment
- Enclosures and expansion units
- Glossary

You can access the POWER9 systems information at [IBM Knowledge Center](#).

Product documentation is also available on DVD (SK5T-7087).

The following information is shipped with the 9040-MR9:

- Power Hardware Information DVD (SK5T-7087)
- Important Notices
- Warranty Information
- License Agreement for Machine Code

Hardware documentation such as installation instructions, user's information, and service information is available to download or view at the [IBM Support Portal](#) website.

The IBM Systems Information Center provides you with a single information center where you can access product documentation for IBM systems hardware, operating systems, and server software. Through a consistent framework, you can efficiently find information and personalize your access. Go the IBM Systems Information Center, at [IBM Knowledge Center](#).

To access the IBM Publications Center Portal, go to the [IBM Publications Center](#) website.

The Publications Center is a worldwide central repository for IBM product publications and marketing material with a catalog of 70,000 items. Extensive search facilities are provided. A large number of publications are available online in various file formats, which can currently be downloaded.

National language support

Not applicable.

Services

IBM Systems Lab Services

IBM Systems Lab Services offers a wide array of services available for your enterprise. It brings expertise on the latest technologies from the IBM development community and can help with your most difficult technical challenges.

IBM Systems Lab Services exists to help you successfully implement emerging technologies so as to accelerate your return on investment and improve your satisfaction with your IBM systems and solutions. Services examples include initial implementation, integration, migration, and skills transfer on IBM systems solution capabilities and recommended practices. IBM Systems Lab Services is one of the service organizations of IBM's world-renowned IBM Systems Group development labs.

For details on available services, contact your IBM representative or go to the [IBM Systems Lab Services](#) website.

Power to Cloud Reward Program

The IBM Power to Cloud Reward Program is designed to accelerate the transformation of your IT infrastructure to private and hybrid cloud, by helping you design, build and deliver a cloud platform on IBM Power Systems servers with help

from IBM Systems Lab Services. You can earn reward points on purchases of select IBM Power Systems servers. Reward points can be used for a range of services focused on helping the transition from traditional IT platforms to private and hybrid cloud platforms by leveraging the proven expertise of IBM Systems Lab Services consultants.

The Power to Cloud Reward option is included with new Power E950 systems and MES upgrades into a Power E950. The Lab Services Power to Cloud team will contact the account team/customer after the shipment to help with the selection of a Power to Cloud service offer. Optionally, the account team or the client may contact the Worldwide Power to Cloud team at pwrcld@us.ibm.com. The customer has up to 90 days from the installation date of the Power E950 to select a Power to Cloud offering. Delivery of the selected service must be completed within 12 months of the installation date of the Power E950 system.

For more details on available Power to Cloud options, see the [IBM Power to Cloud Reward Program](#) website.

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 go to the [IBM Global Technology Services^{\(R\)}](#) website.

For details on available IBM Business Continuity and Recovery Services, contact your IBM representative or go to the [Business Resiliency Services](#) website.

Details on education offerings related to specific products can be found on the [IBM authorized training](#) website.

Technical information

Specified operating environment

Physical specifications

Power E950 (9040-MR9)

System node

- Width: 448 mm (17.5 in.)
- Depth: 902 mm (35.5 in.)
- Height: 175 mm (6.9 in.), 4 EIA units
- Weight: 69 kg (152 lb)

Note: The EMX0 Remote I/O Drawer connection in the T42 and S42 racks will stop the rear door from closing, thus requiring the mandatory 8 inches rack extensions.

PCIe I/O Expansion Drawer

- Width: 448 mm (17.5 in.)
- Depth: 736.6 mm (29 in.)
- Height: 177.8 mm (7.0 in.), 4 EIA units
- Weight: 54.4 kg (120 lb)

To assure installability and serviceability in non-IBM industry-standard racks, review the installation planning information for any product-specific installation requirements.

Operating environment

- Temperature:
 - 5°C - 45°C (41°F - 113°F) nonoperating
 - 18°C - 27°C (64°F - 80°F) recommended operating
 - 10°C - 40°C (50°F - 104°F) allowable operating
 - Derate maximum allowable dry-bulb temperature 1°C (1.8°F) per 175 m (574 ft) above 900 m (2,953 ft) up to a maximum allowable elevation of 3050 m (10,000 ft).
- Relative humidity (noncondensing):
 - 5% - 80% RH nonoperating
 - 8% - 80% operating
- Maximum dew point:
 - 27°C (80°F) nonoperating
 - 24°C (75°F) operating
- Operating voltage: 200 - 240 V AC
 - Operating frequency: 50 - 60 Hz +/- 3 Hz
 - Power consumption: 3,850 watts maximum (per system node)
 - Power source loading: 3.9 kVA maximum (per system node)
 - Thermal output: 14,403 Btu/hr maximum (per system node)
 - Maximum altitude: 3,050 m (10,000 ft)
- **Notes:**
 - To calculate the amperage, multiply the kVA by 1000 and divide that number by the operating voltage.
 - Model 9040-MR9 use four power supply units.
- Noise level: (Acoustics A-weighted Upper-Limit Sound Power Levels)
 - One typically configured four 8-core or 12-core, 2 TB memory:
 - 7.4 bels (operating/idle: 25 C, 500 m)
 - One maximum configured four 12-core, 2 TB memory:
 - 8.1 bels (heavy workload, 25 C, 500 m)

The Power E950 server must be installed in a rack with a rear door and side panels for EMC compliance. The native HMC Ethernet ports must use shielded Ethernet cables.

Note: Government regulations, such as those prescribed by OSHA or European Community Directives, may govern noise level exposure in the workplace and may apply to you and your server installation. This IBM system is available with an optional acoustical door feature that can help reduce the noise emitted from this system. The actual sound pressure levels in your installation depend upon a variety of factors, including the number of racks in the installation; the size, materials, and configuration of the room where you designate the racks to be installed; the noise levels from other equipment; the room ambient temperature, and employees' location in relation to the equipment. Further, compliance with such government regulations also depends upon a variety of additional factors, including the duration of employees' exposure and whether employees wear hearing protection. IBM recommends that you consult with qualified experts in this field to determine whether you are in compliance with the applicable regulations.

EMC conformance classification

This equipment is subject to Electromagnetic Compatibility (EMC) regulations and shall comply with the appropriate country EMC regulations before final delivery to the buyer or centers of distribution.

- US: FCC CFR, Title 47, Part 15, EMI Class A
- EEA, Turkey: EU Council Directive 2014/30/EU, EMI Class A
- Japan: VCCI Council, EMI Class A
- Korea: RRA, EMI Class A
- China (PRC): CPCS, EMI Class A
- Taiwan R.O.C.: BSMI CNS 13438, EMI Class A
- Australia\New Zealand: ACMA, EMI Class A
- Canada: ICES-003, EMI Class A
- Eurasian Economic Area (EAEU), EMI Class A
- Saudi Arabia: MoCI, EMI Class A
- Vietnam: MIC, EMI Class A
- Morocco EMC Order, EMC Class A

Homologation - Telecom type approval

Homologation approval for specific countries has been initiated with the IBM Homologation and Type Approval (HT&A) organization in Nice, France.

The Power System E950 server nodes or system control unit or PCIe Expansion units are not certified for connection to interfaces of public telecommunications networks. Certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions and for information on PCIe adapters that can be used in the system and which are certified.

Product safety/country testing/certification

- UL 60950-1:2007 Underwriters Laboratory
- CAN/CSA22.2 No. 60950-1-07
- EN60950-1:2006 European Norm
- IEC 60950-1 2nd Edition + all National Differences

Homologation

The Power System E950 server nodes or system control unit or PCIe Expansion units are not certified for direct connection to interfaces of public telecommunications networks. Certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions and for information on PCIe adapters that can be used in the system and which are certified.

Software requirements

If installing the AIX operating system LPAR with any I/O configuration (one of these):

- AIX Version 7.2 with the 7200-02 Technology Level and Service Pack 7200-02-02-1832, or later
- AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-02-1832, or later
- AIX Version 7.2 with the 7200-03 Technology Level, or later (planned availability September 14, 2018)
- AIX Version 6.1 with the 6100-09 Technology Level and Service Pack 6100-09-12-1838, or later (planned availability September 14, 2018)(AIX 6.1 service extension required)

- AIX Version 7.2 with the 7200-01 Technology Level and Service Pack 7200-01-05-1845, or later (planned availability January 31, 2019)
- AIX Version 7.1 with the 7100-04 Technology Level and Service Pack 7100-04-07-1845, or later (planned availability January 31, 2019)

If installing the AIX operating system Virtual I/O only LPAR (one of these):

- AIX Version 7.2 with the 7200-03 Technology Level, or later (planned availability September 14, 2018)
- AIX Version 7.2 with the 7200-02 Technology Level and Service Pack 7200-02-01-1732, or later
- AIX Version 7.2 with the 7200-01 Technology Level and Service Pack 7200-01-01-1642, or later
- AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-01-1731, or later
- AIX Version 7.1 with the 7100-04 Technology Level and Service Pack 7100-04-02-1614, or later
- AIX Version 6.1 with the 6100-09 Technology Level and Service Pack 6100-09-07-1614, or later (AIX 6.1 service extension required)

See the [Power Systems Prerequisites](#) website for compatibility information for hardware features and the corresponding AIX and IBM i Technology Levels.

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

- Red Hat Enterprise Linux 7.5 for Power LE (p8compat), or later
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 7 for Power LE version 7.5, or later
- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 11 Service Pack 4, or later
- SUSE Linux Enterprise Server 12 Service Pack 3, or later
- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 12 Service Pack 3, or later
- SUSE Linux Enterprise Server 15, or later

Note: The above list applies to the installation of the Linux operating system LPAR in nonproduction SAP HANA implementations. See the [Statement of general direction](#) for IBM Power System E950 (9040-MR9) for SAP HANA Production Use.

If installing VIOS:

- VIOS 2.2.6.23, or later

Java™ is supported on POWER9 servers. For best use of the performance capabilities and the most recent improvements of POWER9 technology, IBM recommends upgrading Java-based applications to Java 7, Java 8, or later, whenever possible. For those clients who want to run Java in AIX environments, see the [AIX Download and service information](#) website.

For Linux (including PowerLinux™), see the [Linux Download information](#) website.

Recommendation: Clients are strongly encouraged to move to a more current supported version: Java 7, Java 7.1, Java 8, or later.

Limitations

- The PCIe I/O Expansion Drawer (#EMX0) will be limited to a maximum of two for August 7, 2018, announcement.
- IBM Power System 9040-MR9 Rack does not support earthquake regions.

Planning information

Cable orders

No additional cables are required.

Security, auditability, and control

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

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

IBM Systems Lab Services

For details on available services, contact your IBM representative or go to the [IBM Systems Lab Services](#) website.

Global Technology Services

Contact your IBM representative for the list of selected services available in your country, either as standard or customized offerings, for the efficient installation, implementation, or integration of this product.

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, go to the [IBM Electronic Service Agent](#) website.

Terms and conditions

Volume orders

Contact your IBM representative.

Products - terms and conditions

Warranty period

Warranty and Additional Coverage Options:	Coverage summary:
Warranty period:	3 years (1 year of standard warranty + 2 additional years of extended support and warranty services upgrade) ¹
Service Level:	IBM On-Site Repair, 24x7 Same Day
Service Upgrade options :	
Warranty Service Upgrade	IBM On-Site Repair, Same Day, 2 hour response ²
Maintenance Services (Post-Warranty):	IBM On-Site Repair, Same Day
IBM Hardware Maintenance Services - committed maintenance ³ :	Y

¹ Known exception:

China:

Warranty period: 1 Year

Service Level: 9x5 Next Business Day

² Offered in US only

³ Not offered in the US.

To obtain copies of the IBM Statement of Limited Warranty, contact your reseller or IBM. An IBM part or feature installed during the initial installation of an IBM machine is subject to the full warranty period specified by IBM. An IBM part or feature that replaces a previously installed part or feature assumes the remainder of the warranty period for the replaced part or feature. An IBM part or feature added to a machine without replacing a previously installed part or feature is subject to a full warranty. Unless specified otherwise, the warranty period, type of warranty service, and service level of a part or feature are the same as those for the machine in which it is installed.

Any IBM mainstream devices (previously called *read intensive device*) identified in this document have a maximum number of write cycles. IBM mainstream device failures will be replaced during standard warranty and maintenance period for devices that have not reached the maximum number of write cycles. Devices that reach this limit may fail to operate according to specifications and must be replaced at the client's expense. Individual service life may vary and can be monitored using an operating system command.

Extended Warranty Service

This product is provided with one year of standard warranty⁴.

For your convenience, IBM has provided two additional years of extended warranty services, plus an upgrade to 24x7 coverage to make this offering. Consult with your advisers about the appropriate financial treatment for this offering.

⁴ Clients have a choice whether or not to accept the additional two years of extended warranty services plus the upgrade that IBM has provided for this offering. If the choice is to not accept the stated extended warranty services and/or the upgrade, contact your adviser about pricing and other options.

Warranty Services

If required, IBM provides repair or exchange service depending on the types of warranty service specified for the machine. IBM will attempt to resolve your problem over the telephone, or electronically through an IBM website. Certain machines contain remote support capabilities for direct problem reporting, remote problem determination, and resolution with IBM. You must follow the problem determination and resolution procedures that IBM specifies. Following problem determination, if IBM determines on-site service is required, scheduling of service will depend upon the time of your call, machine technology and redundancy, and availability of parts.

If applicable to your product, parts considered Customer Replaceable Units (CRUs) will be provided as part of the machine's standard warranty service.

Service levels are response-time objectives and are not guaranteed. The specified level of warranty service may not be available in all worldwide locations. Additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country-specific and location-specific information.

CRU Service

IBM provides replacement CRUs to you for you to install. CRU information and replacement instructions are shipped with your machine and are available from IBM upon your request. CRUs are designated as being either a Tier 1 (mandatory) or a Tier 2 (optional) CRU.

Tier 1 (mandatory) CRU

Installation of Tier 1 CRUs, as specified in this announcement, is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation.

The following parts have been designated as Tier 1 CRUs:

- Keyboard
- Mouse
- Display
- Fan
- Disk Drive
- RDX drive
- NVMe carrier card
- NVMe flash module
- Op Panel - Base
- Op Panel - LCD
- TOD Battery
- IO Cassette
- PCIe adapter
- Power Supply
- FSP Card
- Memory Riser
- Memory DIMM
- Memory VRM
- IO VRM
- Processor VRM
- TPM Card
- Front Bezel
- Power Cable
- Service Cover

Tier 2 (optional) CRU

You may install a Tier 2 CRU yourself or request IBM to install it, at no additional charge.

Based upon availability, CRUs will be shipped for next-business-day (NBD) delivery. IBM specifies, in the materials shipped with a replacement CRU, whether a defective CRU must be returned to IBM. When return is required, return instructions and a container are shipped with the replacement CRU. You may be charged for the

replacement CRU if IBM does not receive the defective CRU within 15 days of your receipt of the replacement.

The following parts have been designated as Tier 2 CRUs:

Processor module

CRU and On-site Service

At IBM's discretion, you will receive specified CRU service, or IBM will repair the failing machine at your location and verify its operation. You must provide a suitable working area to allow disassembly and reassembly of the IBM machine. The area must be clean, well-lit, and suitable for the purpose.

Service level is:

- 24 hours per day, 7 days a week, , same-day response. Same day service level includes the installation of Tier 1 CRUs at no additional charge

Warranty services

IBM is now shipping machines with selected non-IBM parts that contain an IBM field replaceable unit (FRU) part number label. These parts are to be serviced during the IBM machine warranty period. IBM is covering the service on these selected non-IBM parts as an accommodation to their customers, and normal warranty service procedures for the IBM machine apply.

International Warranty Service

International Warranty Service allows you to relocate any machine that is eligible for International Warranty Service and receive continued warranty service in any country where the IBM machine is serviced. If you move your machine to a different country, you are required to report the machine information to your Business Partner or IBM representative.

The warranty service type and the service level provided in the servicing country may be different from that provided in the country in which the machine was purchased. Warranty service will be provided with the prevailing warranty service type and service level available for the eligible machine type in the servicing country, and the warranty period observed will be that of the country in which the machine was purchased.

The following types of information can be found on the [International Warranty Service](#) website

- Machine warranty entitlement and eligibility
- Directory of contacts by country with technical support contact information
- Announcement Letters

Warranty service upgrades

During the warranty period, warranty service upgrades provide an enhanced level of On-site Service for an additional charge. Service levels are response-time objectives and are not guaranteed. See the [Warranty services](#) section for additional details.

IBM will attempt to resolve your problem over the telephone or electronically by access to an IBM website. Certain machines contain remote support capabilities for direct problem reporting, remote problem determination, and resolution with IBM. You must follow the problem determination and resolution procedures that IBM specifies. Following problem determination, if IBM determines on-site service is required, scheduling of service will depend upon the time of your call, machine technology and redundancy, and availability of parts.

CRU and On-site Service

At IBM's discretion, you will receive CRU service or IBM will repair the failing machine at your location and verify its operation. You must provide a suitable working area to allow disassembly and reassembly of the IBM machine. The area must be clean, well-lit, and suitable for the purpose.

Customer Replaceable Units (CRUs) may be provided as part of the machine's standard warranty CRU Service except that you may install a CRU yourself or request IBM installation, at no additional charge, under the CRU and On-site Service level specified above. For additional information on the CRU Service, see the warranty information.

Maintenance services

If required, IBM provides repair or exchange service depending on the types of maintenance service specified for the machine. IBM will attempt to resolve your problem over the telephone or electronically, through an IBM website. Certain machines contain remote support capabilities for direct problem reporting, remote problem determination, and resolution with IBM. You must follow the problem determination and resolution procedures that IBM specifies. Following problem determination, if IBM determines on-site service is required, scheduling of service will depend upon the time of your call, machine technology and redundancy, and availability of parts. Service levels are response-time objectives and are not guaranteed. The specified level of maintenance service may not be available in all worldwide locations. Additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country-specific and location-specific information.

The following service selections are available as maintenance options, at additional cost, for your machine type.

- IBM On-Site Repair, Next Business Day Response Target, Monday-Friday. Calls must be received by 3:00 PM local time in order to qualify for next business day response.
- IBM On-Site Repair, Same Day On-Site Response Target, Monday-Sunday 00:00-24:00, 365 days a year.⁵

⁵ Recommended for this product

IBM Hardware Maintenance Services - committed maintenance⁶

Organizations can lose as much as USD \$100 million per year to downtime related to information and communications technology. IBM Hardware Maintenance Services -- committed maintenance can deliver different types of hardware service for IBM equipment from the moment you call for support worldwide (based on the countries in which IBM has a presence) and around the clock. Through clear response targets and standardized service-delivery metrics, we can help you optimize your IT infrastructure and can help reduce the threat of hardware-related outages.

For more information, see the [IBM Hardware Maintenance Services - committed maintenance Data Sheet \(D\) - USEN](#).

⁶ Currently not available in the US.

On-site Service

IBM will repair the failing machine at your location and verify its operation. You must provide a suitable working area to allow disassembly and reassembly of the IBM machine. The area must be clean, well-lit, and suitable for the purpose.

Customer Replaceable Unit (CRU) Service

If your problem can be resolved with a CRU (for example, keyboard, mouse, speaker, memory, or hard disk drive), and depending upon the maintenance service offerings in your geography, IBM will ship the replacement CRU to you for you to

install. CRU information and replacement instructions are shipped with your machine and are available from IBM upon your request.

CRUs will be shipped based upon availability. IBM specifies, in the materials shipped with a replacement CRU, whether a defective CRU must be returned to IBM. When return is required, 1) return instructions and a container are shipped with the replacement CRU, and 2) you may be charged for the replacement CRU if IBM does not receive the defective CRU within 15 days of your receipt of the replacement.

CRUs are designated as being either a Tier 1 (mandatory) or a Tier 2 (optional) CRU.

Tier 1 (mandatory) CRUs: Installation of Tier 1 CRUs, as specified in this announcement, is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation.

For machines with On-site Same-day Response Service, IBM will replace a Tier 1 CRU part at your request, at no additional charge.

The following parts have been designated as Tier 1 CRUs:

- Keyboard
- Mouse
- Display
- Fan
- Disk Drive
- RDX drive
- NVMe carrier card
- NVMe flash module
- Op Panel Base
- Op Panel LCD
- TOD Battery
- IO Cassette
- PCIe adapter
- Power Supply
- FSP Card
- Memory Riser
- Memory DIMM
- Memory VRM
- IO VRM
- Processor VRM
- TPM Card
- Front Bezel
- Power Cable
- Service Cover

Tier 2 (optional) CRUs: You may install a Tier 2 CRU yourself or request IBM to install it, at no additional charge.

Feature codes or models for which there is a maintenance charge:

9040-MR9

Additional reference for Europe

The following European documents can be found on the [IBM Maintenance and Technical Support Services](#) website.

- European Announcement Letter ZS03-0150 for IBM Customer Agreement (ICA)
- European Announcement Letter ZS04-0135 for Enterprise Agreement Contract
- European Announcement Letter ZS98-0118 for ServiceSuite Contract

Non-IBM parts service

Under certain conditions, IBM provides services for selected non-IBM parts at no additional charge for machines that are covered under warranty service upgrades or maintenance services.

This service includes hardware problem determination (PD) on the non-IBM parts (for example, adapter cards, PCMCIA cards, disk drives, memory) installed within IBM machines and provides the labor to replace the failing parts at no additional charge.

If IBM has a Technical Service Agreement with the manufacturer of the failing part, or if the failing part is an accommodations part (a part with an IBM FRU label), IBM may also source and replace the failing part at no additional charge. For all other non-IBM parts, customers are responsible for sourcing the parts. Installation labor is provided at no additional charge, if the machine is covered under a warranty service upgrade or a maintenance service.

Usage plan machine

No

IBM hourly service rate classification

Two

When a type of service involves the exchange of a machine part, the replacement may not be new, but will be in good working order.

Maintenance service offerings

This machine is eligible under terms and conditions of IBM ServiceElite, the IBM Enterprise Service Agreement (ESA), or the IBM Maintenance Agreement. Consult your IBM representative for details.

General terms and conditions

Field-installable features

Yes

Model conversions

Yes

Machine installation

Customer setup. Customers are responsible for installation according to the instructions IBM provides with the machine.

Graduated program license charges apply

No

Licensed Machine Code

IBM Machine Code is licensed for use by a customer on the IBM machine for which it was provided by IBM under the terms and conditions of the IBM License Agreement for Machine Code, to enable the machine to function in accordance with its specifications, and only for the capacity authorized by IBM and acquired by the customer. You can obtain the agreement by contacting your IBM representative. It

can also be found on the [License Agreement for Machine Code and Licensed Internal Code](#).

Access to Machine Code updates is conditioned on entitlement and license validation in accordance with IBM policy and practice. IBM may verify entitlement through customer number, serial number, electronic restrictions, or any other means or methods employed by IBM in its discretion.

If the machine does not function as warranted and your problem can be resolved through your application of downloadable Machine Code, you are responsible for downloading and installing these designated Machine Code changes as IBM specifies. If you would prefer, you may request IBM to install downloadable Machine Code changes; however, you may be charged for that service.

Machine Code License Acceptance Requirement

B.) Acceptance-By-Use Machine: No, the Machine Code license requires signed acceptance by the machine's end user directly with IBM, applicable to orders for a new machine, machine type conversion MES, and to machines transferred to another user.

Prices

For all local charges, contact your IBM representative.

IBM Global Financing

IBM Global Financing offers competitive financing to credit-qualified customers to assist them in acquiring IT solutions. Offerings include financing for IT acquisition, including hardware, software, and services, from both IBM and other manufacturers or vendors. Offerings (for all customer segments: small, medium, and large enterprise), rates, terms, and availability can vary by country. Contact your local IBM Global Financing organization or go to the [IBM Global Financing](#) website for more information.

IBM Global Financing offerings are provided through IBM Credit LLC in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government customers. Rates are based on a customer's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension, or withdrawal without notice.

Announcement countries

All European, Middle Eastern, and African countries, except Islamic Republic of Iran, Sudan, and Syrian Arab Republic.

Trademarks

Power Systems, POWER9, Active Memory, IBM Cloud, PowerLinux, HACMP, PowerVP and Electronic Service Agent are trademarks of IBM Corporation in the United States, other countries, or both.

IBM, Power, AIX, POWER8, PowerVM, POWER, PowerHA, IBM API Connect, WebSphere and Global Technology Services are registered trademarks of IBM Corporation in the United States, other countries, or both.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis.

Oracle and Java are trademarks of Oracle and/or its affiliates in the United States, other countries, or both.

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

Terms of use

IBM products and services which are announced and available in your country can be ordered under the applicable standard agreements, terms, conditions, and prices in effect at the time. IBM reserves the right to modify or withdraw this announcement at any time without notice. This announcement is provided for your information only. Reference to other products in this announcement does not necessarily imply those products are announced, or intend to be announced, in your country. Additional terms of use are located at

[Terms of use](#)

For the most current information regarding IBM products, consult your IBM representative or reseller, or go to the IBM worldwide contacts page

[IBM Directory of worldwide contacts](#)

Corrections

(Corrected on April 1, 2021)

Product number and Prices sections were revised.

(Corrected on November 21, 2019)

Changed length of CMC service in the Overview section.

(Corrected on August 9, 2019)

Added IBM Proactive Support for Power Systems topic to Overview and Description sections.

(Corrected on May 16, 2019)

The Description, Terms and conditions, and Product number sections were revised.

(Corrected on August 29, 2018)

Revised the Description, Product structure - Ordering information, and Terms and conditions sections.

(Corrected on August 22, 2018)

Revised the Planned availability date section.