

Power E1080 Enterprise server delivers a uniquely architected platform to help securely and efficiently scale core operational and AI applications in a hybrid cloud

Table of contents

1 Overview	35 Publications
4 Key requirements	36 Technical information
4 Planned availability date	39 Terms and conditions
5 Description	45 Prices
23 Statement of general direction	46 Announcement countries
23 Product number	46 Corrections

At a glance

The IBM^(R) Power^(R) E1080, the most powerful and scalable server in the IBM Power portfolio, provides the following underlying hardware components:

- Up to 240 Power10 technology-based processor cores
- Up to 64 TB memory
- Up to 32 Peripheral Component Interconnect^(R) Express (PCIe) Gen5 slots in system nodes
- Up to 192 PCIe Gen3 slots with expansion drawers
- Up to over 4,000 directly attached serial-attached SCSI (SAS) disks or solid-state drives (SSDs)
- Up to 1,000 virtual machines (VMs) per system
- System control unit, providing redundant system master Flexible Service Processor (FSP)

The Power E1080 supports:

- IBM AIX^(R), IBM i, and Linux^(R) environments
- Capacity on demand (CoD) processor and memory options
- IBM Power System Private Cloud Solution with Dynamic Capacity

Overview

Power E1080

Clients need applications and data to be enterprise-grade everywhere without adding complexity and cost. The Power E1080 is the newest addition to IBM Power, the industry's best-in-class server platform for security and reliability. The Power E1080 introduces the essential enterprise hybrid cloud platform--uniquely architected to help you securely and efficiently scale core operational and AI applications in a hybrid cloud. The Power E1080 simplifies end-to-end encryption and brings AI where your data resides for faster insights. This helps enable greater workload flexibility and agility while accomplishing more work. The Power E1080 can help you:

- **Respond faster to business demands** with unmatched performance for efficient scaling and consistent pay-for-use consumption across public and private clouds
- **Protect data with from core to cloud** using full memory encryption at the processor level to support end-to-end security across public and private clouds without impacting performance
- **Streamline insights and automation** by running AI inferencing directly where your operational data resides
- **Maximize availability and reliability** with built-in advanced recovery and self-healing for infrastructure redundancy and disaster recovery in IBM Cloud^(R)

Power E1080 brings AI to where your operational data resides

You can drive business insights faster, meet service level agreements (SLAs), and eliminate security risk associated with data movement by bringing AI to where your data resides.

Each new Power10 processor single chip module (SCM) contains two memory controllers. Four 10-core 3.65 - 3.90 GHz (max), four 12-core 3.6 - 4.15 GHz (max), or four 15-core 3.55 - 4.00 GHz (max) are used in each system node, providing 40 cores to a 160-core system (#EDP2), 48 cores to a 196-core system (#EDP3), or 60 cores to a 240-core system (#EDP4)*. As few as 16 cores in the system can be activated or up to 100% of the cores in the system can be activated. Increments of one core at a time is available through built-in Capacity Upgrade on Demand (CUoD) functions to the full capacity of the system.

The system control unit provides redundant FSPs, the operator panel, and the Vital Product Data (VPD).

An optional external DVD can be attached with a USB cable when a USB adapter is installed in a node.

The memory supported in this server are the next-generation differential dual inline memory modules (DIMMs) implemented by Power called DDIMMs, which utilize DDR4 DRAM memory.

Power E1080 memory options are available as 128 GB (#EMC1), 256 GB (#EMC2), 512 GB (#EMC3), and 1024 GB (#EMC4) memory features. Each memory feature provides four DDIMMs. Each system node supports a maximum of 16 memory features and up to 64 DDIMM slots. Using 1024 GB DDIMM features yields a maximum of 16 TB per node. A two-node system has a maximum of 32 memory features and 32 TB capacity. A four-node system has a maximum of 64 TB capacity. Minimum memory activations of 50% of the installed capacity are required.

The 19-inch PCIe 4U I/O expansion drawer (#EMX0) provides 12 additional slots for PCIe adapters. Up to four PCIe I/O expansion drawers can be attached per system node. For example, a two-node system can have a maximum of eight PCIe I/O expansion drawers for a total of 96 PCIe slots in the I/O drawers with no PCIe slots in the system node.

Direct attached storage is supported with the EXP24SX SFF Gen2-bay drawer (#ESLS), an expansion drawer with 24 2.5-inch form-factor SAS bays.

IBM Power Private Cloud Solution with Dynamic Capacity

The Power Private Cloud Solution with Dynamic Capacity is an infrastructure offering that enables you to take advantage of cloud agility and economics while getting the same business continuity and flexibility that you already enjoy from Power servers. The Power Private Cloud Solution offers:

- Cost optimization with pay-for-use pricing
- Automated and consistent IT management with Red Hat^(R) Ansible^(R) for Power
- IBM Proactive Support for Power systems services
- IBM Systems Lab Services Assessment and implementation assistance

Both Elastic and Shared Utility Capacity options are now available on all Power E1080 systems.

Elastic Capacity on Power E1080 systems enables clients to deploy pay-for-use consumption of processor, memory and supported operating systems, by the day, across a collection of Power E1080 and Power E980 systems.

Shared Utility Capacity on Power E1080 systems provides enhanced multisystem resource sharing and by-the-minute tracking and consumption of compute resources across a collection of systems within a Power Enterprise Pool (2.0). It delivers a complete range of flexibility to tailor initial system configurations with the right mix of purchased and pay-for-use consumption of processor, memory, and software. Clients with existing Power Enterprise Pools of Power E980 systems can simply add Power E1080 systems into their pool and migrate to them at the rate and pace of their choosing, as any Power E980 and Power E1080 server may seamlessly interoperate and share compute resources within the same pool.

A Power Private Cloud Solution infrastructure consolidated onto Power E1080 systems has the potential to greatly simplify system management so IT teams can focus on optimizing their business results instead of moving resources around within their data center.

Shared Utility Capacity resources are easily tracked by virtual machine (VM) and monitored by an IBM Cloud Management Console (CMC), which integrates with local Hardware Management Consoles (HMC) to manage the pool and track resource use by system and VM, by the minute, across a pool.

You no longer need to worry about overprovisioning capacity on each system to support growth, as all available processor and memory on all systems in a pool are activated and available for use.

Base Capacity for processor, memory, and supported operating system entitlement resources is purchased on each Power E980 and Power E1080 system and is then aggregated across a defined pool of systems for consumption monitoring.

Metered Capacity is the additional installed processor and memory resource above each system's Base Capacity. It is activated and made available for immediate use when a pool is started, then monitored by the minute by a CMC.

Metered resource usage is charged only for minutes exceeding the pool's aggregate Base resources, and usage charges are debited in real-time against your purchased Capacity Credits (5819-CRD) on account.

IBM offers a Private Cloud Capacity Assessment and Implementation Service performed by Systems Lab Services professionals, which can be preselected at time of purchase or requested for qualifying Power E1080 servers.

Power to Cloud services

To assist clients with their move to the cloud, IBM is bundling 10,000 points with every Power E1080 server purchase that can be redeemed for onsite cloud deployment services. For additional details, see the [IBM Power to Cloud Rewards Program](#) website. For those clients looking to create their own private cloud, expert services are available for cloud provisioning and automation with IBM Cloud PowerVC Manager with a heavy focus on creating and supporting a DevOps cloud implementation.

For those clients looking for a hybrid cloud solution, Design for Hybrid Cloud Workshop services are available to help you produce best-of-breed applications using IBM API Connect^(R) and IBM Cloud with IBM Power.

To learn more about all the new cloud capabilities that come with the Power E1080 server, see the [IBM Power Enterprise Cloud Index](#) website.

CMC for Power

The CMC is a cloud-native platform that provides apps that give powerful insights into your Power infrastructure across data centers and geographies. With no additional software or infrastructure setup, you can get single pane of glass views of your inventory, software levels, and resource capacity utilization, as well as launch-in-context of your on-premises software, such as IBM PowerVC and IBM PowerHA[®].

Power E1080 server Power10 hardware components

- Ten, twelve, or fifteen core processors
- Up to 240 Power10 processor cores in one to four systems nodes; up to 64 TB of 2933 MHz, DDR4 DRAM memory, and six PCIe Gen4 x16 or PCIe Gen5 x8 and two PCIe Gen5 x8 I/O expansion slots per system node enclosure, with a maximum of 32 per system
- Redundant clocking in each system node
- Four non-volatile memory express (NVMe) drive bays per system node for boot purposes
- System control unit, providing redundant system master FSP and support for the operations panel, the system VPD, and external attached DVD
- 19-inch PCIe Gen3 4U I/O expansion drawer and PCIe fan-out modules, supporting a maximum of 192 PCIe slots and four I/O expansion drawers per node
- PCIe Gen1, Gen2, Gen3, Gen4, and Gen5 adapter cards supported in the system node, and PCIe Gen1, Gen2, Gen3, and Gen4 adapter cards supported in I/O expansion drawer
- EXP24SX SFF drawer with 24 2.5-inch form-factor SAS bays
- Dynamic LPAR support for adjusting workload placement of processor and memory resources
- CoD for processors and memory to help respond more rapidly and seamlessly to changing business requirements and growth
- Active Memory Expansion (AME) that is optimized onto the processor chip
- Active Memory Mirroring (AMM) to enhance resilience by mirroring critical memory used by the PowerVM[®] hypervisor
- Power Enterprise Pools that support unsurpassed enterprise flexibility for workload balancing and system maintenance

*EDP4 is not available to order in China.

Key requirements

The Power E1080 server requires an AIX, Linux, or IBM i operating system. See the [Software requirements](#) section for detailed requirements.

Planned availability date

September 17, 2021, with the following limitations per system:

- EDN1 Initial max = 2
- EMX0 Initial/MES/Both max = 8
- EMXH Initial/MES/Both max = 16
- EJ24 Initial/MES/Both max = 16

October 21, 2021, for South Korea

December 10, 2021, for the following features per system:

- EDN1 Initial max = 4

- EMX0 Initial/MES/Both max = 16
- EMXH Initial/MES/Both max = 32
- EJ24 Initial/MES/Both max = 32
- EMC3 and EMC4
- EFCF and EFCG
- EN1F and EN1H

Note: The following features are not orderable in China: EDP4, EDPD, EDAD, EPDS, EPDX, EMEG, EMEL, EPS2, EPS7, EPSM, ELCM, and EME4.

Description

Security, operational efficiency, and real-time intelligence to respond quickly to market changes are now nonnegotiable for IT. In an always-on environment of constant change, you need to automate and accelerate critical operations, while ensuring 24/7 availability and staying ahead of cyberthreats. You need applications and data to be enterprise-grade everywhere, but without adding complexity and cost.

Power servers are already the most reliable and secure in their class. Today, the new Power E1080 extends that leadership and introduces the essential enterprise hybrid cloud platform--uniquely architected to help you securely and efficiently scale core operational and AI applications anywhere in a hybrid cloud. Now you can encrypt all data simply without management overhead or performance impact and drive insights faster with AI at the point of data. You can also gain workload deployment flexibility and agility with a single hybrid cloud currency while doing more work.

Power E1080 feature summary

The following features are available on the Power E1080 server:

- One to four 5U system nodes
 - The Power E1080 server will support three and four systems nodes by December 10, 2021.
- One 2U system control unit
- One to four processor features per system with four single-chip modules (SCMs) per feature:
 - 3.65 - 3.90 GHz, 40-core Power10 processor (#EDP2)
 - 3.6 - 4.15 GHz, 48-core Power10 processor (#EDP3)
 - 3.55 - 4.00 GHz, 60-core Power10 processor (#EDP4)*
- CoD processor core activation features available on a per-core basis
- 64 DDIMM slots per system node
- DDR4 DDIMM memory cards:
 - 128 GB (4 x 32 GB), (#EMC1)
 - 256 GB (4 x 64 GB), (#EMC2)
 - 512 GB (4 x 128 GB), (#EMC3)
 - 1024 GB (4 x 256 GB), (#EMC4)
- CoD memory activation features include:
 - 100 GB Mobile Memory Activations (#EDAB)
 - 500 GB Mobile Memory Activations (#EMBK)
- AME optimized onto the processor chip (#EM8F)
- Six PCIe Gen4 x16 or PCIe Gen5 x8 and two PCIe Gen5 x8 I/O low-profile expansion slots per system node (maximum 32 in a 4-node system)
- One USB port to support external attached DVD when a USB adapter is installed in a node

- Redundant hot-swap AC power supplies in each system node drawer
- Two HMC ports in the system control unit
- Optional PCIe I/O expansion drawer with PCIe slots:
 - Zero to four drawers per system node drawer (#EMX0).
 - Each I/O drawer holds one or two 6-slot PCIe fan-out modules (#EMXH).
 - Each fan-out module attaches to the system node through a PCIe optical cable adapter (#EJ24).

System nodes

Each 5 EIA or 5U system node of the server has four air-cooled SCMs optimized for performance and scalability. The Power E1080 SCMs can have ten, twelve, and fifteen Power10 cores running at up to 4.15 GHz and simultaneous multithreading that executes up to eight threads per core. Each SCM has dual memory controllers to deliver up to 409 GBps of peak memory bandwidth per socket or 1636 GBps per node. Using PCIe Gen5 I/O controllers, which are also integrated onto each SCM to further reduce latency, up to 576 GBps peak I/O bandwidth is available per node. Thus, a Power E1080 system bandwidth can help provide maximum processor performance, enabling applications to run faster and be more responsive.

Each system node has 64 DDIMM slots and can support up to 16 TB of DDR4 memory. Thus, a four-node server can have up to 64 TB of memory. The system node has four internal NVMe U.2 (2.5-in. 7mm form factor) SSDs. Each SSD is driven from a x4 PCIe Gen4 connection. Each system node has eight PCIe slots, which six are Gen4 x16 or Gen5 x8 and two are Gen5 x8, low profile. Thus, a four-node server can have up to 32 PCIe slots. PCIe expansion units can optionally expand the number of PCIe slots on the server.

A system node is ordered using a processor feature. Each processor feature will deliver a set of four identical SCMs in one system node. All processor features in the system must be identical. Cable features are required to connect system node drawers to the system control unit and to other system nodes.

Processor core activations

Each Power E1080 server requires a minimum of 16 permanent processor core activations, using either static activations or Linux on Power activations. This minimum is per system, not per node. The rest of the cores can be permanently or temporarily activated or remain inactive until needed. The activations are not specific to hardware cores, SCMs, or nodes. They are known to the system as a total number of activations of different types and used or assigned by the Power hypervisor appropriately.

A variety of activations fit different usage and pricing options. Static activations are permanent and support any type of application environment on this server. Mobile activations are ordered against a specific server, but they can be moved to any server within the Power Enterprise Pool and can support any type of application.

60-core (#EDP4)*	48-core (#EDP3)	40-core (#EDP2)
1-core static activation (#EDPD)*	1-core static activation (#EDPC)	1-core static activation (#EDPB)
1-core Power Linux (#ELCM)*	1-core Power Linux (#ELCQ)	1-core Power Linux (#ELCL)

Memory

Differential DIMMs (DDIMMs) are extremely high-performance, high-reliability, intelligent, and DRAM memory. DDR4 technology is employed and provide the performance. DDIMMs are placed in DDIMM slots in the system node.

Each system node has 64 memory DDIMM slots, and at least half of the memory slots are always physically filled. Sixteen DDIMM slots are local to each of the four SCMs in the server, but SCMs and their cores have access to all the other memory in the server. When filling the other memory slots in the SCM, a quantity of four DDIMMs must be used. Thus, the DDIMM slots of the SCMs are from 50% to 100%

filled. The system node (four SCMs) DDIMM slots can have 32, 36, 40, 44, 48, 52, 56, 60, and 64 DDIMMs physically installed (quad plugging rules).

To assist with the quad plugging rules, four DDIMMs are ordered using one memory feature number. Select from 128 GB feature EMC1 (4 x 32 DDR4), 256 GB feature EMC2 (4 x 64 DDR4), 512 GB feature EMC3 (4 x 128 DDR4), or 1024 GB feature EMC4 (4 x 256 DDR4).

All DDIMMs must be identical on the same SCM, so if you're using eight DDIMMs, both memory features on an SCM must be identical. A different SCM in the same system node can use a different memory feature. For example, one system node could technically use 128 GB, 256 GB, 512 GB, and 1024 GB memory features.

To provide more flexible pricing, memory activations are ordered separately from the physical memory and can be permanent or temporary. Activation features can be used on DDR4 memory features and used on any size memory feature. Activations are not specific to a DDIMM, but they are known as a total quantity to the server. The Power hypervisor determines what physical memory to use.

Memory activation features are:

- 1 GB Memory Activations (#EMAZ), static
- 100 GB Memory Activations (#EMQZ), static
- 100 GB Mobile Memory Activations (#EDAB)
- 500 GB Memory Activations for Power Linux (#ELME)

A minimum of 50% of the total physical memory capacity of a server must have permanent memory activations ordered for that server. For example, a server with a total of 8 TB of physical memory must have at least 4 TB of permanent memory activations ordered for that server. These activations can be static, mobile, or Linux on Power. At least 25% must be static activations or Linux on Power activations. For example, a server with a total of 8 TB physical memory must have at least 2 TB of static activations or Linux on Power activations. The 50% minimum cannot be fulfilled using mobile activations ordered on a different server.

The minimum activations ordered with MES orders of additional physical memory features will depend on the existing total installed physical memory capacity and the existing total installed memory activation features. If you already have installed more than 50% activations for your existing system, then you can order less than 50% activations for the MES ordered memory. The resulting configuration after the MES order of physical memory and any MES activations must meet the same 50% and 25% minimum rules above.

For the best possible performance, it is generally recommended that memory be installed evenly across all system node drawers and all SCMs in the system. Balancing memory across the installed system planar cards enables memory access in a consistent manner and typically results in better performance for your configuration.

Though maximum memory bandwidth is achieved by filling all the memory slots, plans for future memory additions should be considered when deciding which memory feature size to use at the time of initial system order.

The AME is an option that can increase the effective memory capacity of the system. See the AME information later in this section.

Power Enterprise Pools with Mobile and Shared Utility Capacity

Power Enterprise Pools 2.0 is the strategic, automated resource sharing technology designed for Power E980 and E1080 systems.

The following capabilities are designed to provide a smooth migration path from Power E980 systems to Power E1080 systems within the same Power Enterprise Pool:

- Capacity Credits for Power (5819-CRD) may be applied to a Power Enterprise Pool (2.0) containing a combination of Power E1080 and E980 systems. Metered Capacity resources consumed will be debited at the same rate for both E1080 and E980 systems.
- For each Base Activation feature purchased new on a Power E1080 server that is replacing a Power E980 server in the same Power Enterprise Pool, up to three (3) Base Activation features may be exchanged from the E980 server for three (3) new, corresponding Power E1080 Base Activation features, at no additional charge, when the Power E980 system is being removed from the pool.
- To support clients migrating to Power E1080 from Power E980 systems in a Power Enterprise Pool (1.0) with Mobile Capacity, a one-time migration is being enabled to allow a quantity of POWER9 Mobile Processor and Memory activation features, purchased initially on a Power E980 system, to migrate and convert to similar features on a Power E1080 system, within the same Power Enterprise Pool, at no additional charge, when the Power E980 system is being removed from the pool.
- All offers of exchange of Base Capacity features and migration of Mobile Capacity features are designed to be executed through the [IBM Entitled Systems Support](#) portal, and are subject to its availability within a country.

System control unit

The 2U system control unit provides redundant system master FSP. Additionally, it contains the operator panel and the system VPD. One system control unit is required for each server. A unique feature number is not used to order the system control unit. One is shipped with each Power E1080 server. Two FSPs in the system control unit are ordered using two EDFP features. All system nodes connect to the system control unit using the cable features EFCH, EFCE, EFCF, and EFCG.

The system control unit is powered from the system nodes. UPIC cables provide redundant power to the system control unit. In a single node system, two UPIC cables are attached to system node 1. In a two-node, three-node, or four-node system, one UPIC cable attaches to system node 1 and one UPIC cable attaches to system node 2. They are ordered with features EFCH. Only one UPIC cable is enough to power the system control unit, and the others are in place for redundancy.

System node PCIe slots

- Each system node enclosure provides excellent configuration flexibility and expandability with eight half-length, low-profile (half-high) PCIe Gen5 slots. The slots are labeled C0 through C7. C0, C1, C2, C5, C6, and C7 are x16, and C3 and C4 are x8.
- These PCIe slots can be used for either low-profile PCIe adapters or for attaching a PCIe I/O drawer.
- A blind swap cassette (BSC) is used to house the low-profile adapters that go into these slots. The server is shipped with a full set of BSCs, even if the BSCs are empty. A feature number to order additional low-profile BSCs is not required or announced.
- If additional PCIe slots beyond the system node slots are required, a system node x16 slot is used to attach a six-slot expansion module in the I/O drawer. An I/O drawer holds two expansion modules that are attached to any two x16 PCIe slots in the same system node or in different system nodes.
- PCIe Gen1, Gen2, Gen3, Gen4, and Gen5 adapter cards are supported in these Gen5 slots. The set of PCIe adapters that are supported is found in the Sales Manual, identified by feature number.
- Concurrent repair and add/removal of PCIe adapter cards is done by HMC-guided menus or by operating system support utilities.
- The system nodes sense which PCIe adapters are installed in their PCIe slots; if an adapter requires higher levels of cooling, they automatically speed up the fans to increase airflow across the PCIe adapters.

PCIe I/O expansion drawer

The 19-inch PCIe 4U I/O expansion drawer (#EMX0) provides slots to hold PCIe adapters that cannot be placed into a system node. The PCIe I/O expansion drawer (#EMX0) and two PCIe fan-out modules (#EMXH) provide 12 PCIe I/O full-length, full-height slots. One fan-out module provides six PCIe slots labeled C1 through C6. The C1 and C4 are x16 slots, and C2, C3, C5, and C6 are x8 slots.

PCIe Gen1, Gen2, Gen3, and Gen4 and 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 9080-HEX for details and rules associated with specific adapters supported and their supported placement in x8 or x16 slots.

Up to four PCIe I/O drawers per node can be attached to the Power E1080 server. Using two 6-slot fan-out modules per drawer provides a maximum of 48 PCIe slots per system node. With two system nodes, up to 96 PCIe slots (8 I/O drawers) are supported. With a 4-node Power E1080 server, up to 192 PCIe slots (16 I/O drawers) are supported.

Additional PCIe I/O drawer configuration flexibility is provided to the Power E1080 servers. Zero, one, two, three, or four PCIe I/O drawers can be attached per system node. As an alternative, a half drawer that consists of just one PCIe fan-out module in the I/O drawer is also supported, enabling a lower-cost configuration if fewer PCIe slots are required. Thus, a system node supports the following half-drawer options: one half drawer, two half drawers, three half drawers, or four half drawers. Because there is a maximum of four feature EMX0 drawers per node, a single system node cannot have more than four half drawers. A server with more system nodes can support more half drawers up to four per node. A system can also mix half drawers and full PCIe I/O drawers. The maximum of four PCIe drawers per system node applies whether a full or half PCIe drawer.

PCIe drawers can be concurrently added to the server at a later time. The drawer being added can have either one or two fan-out modules. Note that adding a second fan-out module to a half-full drawer does require scheduled downtime.

PCIe I/O drawer attachment and cabling

- A PCIe x16 to optical CXP converter adapter (#EJ24) and 2.0 M (#ECCR), 10.0 M (#ECCY), or 20 M (#ECCZ) CXP 12X Active Optical Cables (AOC) connect the system node to a PCIe fan-out module in the I/O expansion drawer. One ECCR, ECCY, or ECCZ feature ships two AOC cables from IBM.
- The two AOC cables connect to two CXP ports on the fan-out module and to two CXP ports on the feature EJ24 adapter. The top port of the fan-out module must be cabled to the top port of the feature EJ24 port. Likewise, the bottom two ports must be cabled together.
- It is recommended but not required that one I/O drawer be attached to two different system nodes in the same server (one drawer module attached to one system node and the other drawer module attached to a different system node). This can help provide cabling for higher availability configurations.
- It is generally recommended that any attached PCIe I/O expansion drawer be located in the same rack as the Power10 server for ease of service, but expansion drawers can be installed in separate racks if the application or other rack content requires it. If you are attaching a large number of cables, such as SAS cables or CAT5/CAT6 Ethernet cables, to a PCIe I/O drawer, then it is generally better to place that feature EMX0 drawer in a separate rack for better cable management.

Limitation: When this cable is ordered with a system in a rack specifying IBM Plant integration, IBM Manufacturing will ship SAS cables longer than three meters in a separate box and not attempt to place the cable in the rack. This is because the longer SAS cable is probably used to attach to an EXP24S drawer 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.

- A BSC is used to house the full-high adapters that go into these slots. The BSC is the same BSC as used with 12X attached I/O drawers (#5802, #5803, #5877, #5873) of the previous-generation server. The drawer is shipped with a full set of BSCs, even if the BSCs are empty. A feature to order additional full-high BSCs is not required or announced.
- A maximum of 16 EXP24s drawers are needed per PCIe drawer (#EMX0) to enable SAS cables to be properly handled by the feature EMX0 cable management bracket.

EXP24SX disk/SSD drawer

- Direct attached storage is supported with the EXP24SX SFF Gen2-bay drawer (#ESLS), an expansion drawer with 24 2.5-inch form-factor SAS bays.
- The Power E1080 server supports up to 4,032 drives with a maximum of 168 EXP24SX drawers. The maximum of 16 EXP24SX drawers per PCIe I/O drawer due to cabling considerations remains unchanged.
- The EXP24SX SFF Gen2-bay drawer (#ESLS) is an expansion drawer with 24 2.5-inch form-factor SAS bays. Slot filler panels are included for empty bays when initially shipped. The EXP24SX supports up to 24 hot-swap SFF-2 SAS HDDs or SSDs. It uses only two EIA of space in a 19-inch rack. The EXP24SX includes redundant AC power supplies and uses two power cords.
- With AIX, Linux, and VIOS, you can order the EXP24SX with four sets of 6 bays, two sets of 12 bays, or one set of 24 bays (mode 4, 2, or 1). With IBM i, you can order the EXP24SX as one set of 24 bays (mode 1). Mode setting is done by IBM Manufacturing, and there is no option provided to change the mode after it is shipped from IBM. Note that when changing modes, a skilled, technically qualified person should follow the special documented procedures. Improperly changing modes can potentially destroy existing RAID sets, prevent access to existing data, or allow other partitions to access another partition's existing data. Hire an expert to assist if you are not familiar with this type of reconfiguration work.
- The EXP24SX SAS ports are attached to a SAS PCIe adapter or pair of adapters using SAS YO or X cables.
- To maximize configuration flexibility and space utilization, the system node does not have integrated SAS bays or integrated SAS controllers. PCIe adapters and the EXP24SX can be used to provide direct access storage.
- To further reduce possible single points of failure, EXP24SX configuration rules consistent with previous Power servers are used. IBM i configurations require the drives to be protected (RAID or mirroring). Protecting the drives is highly recommended, but not required for other operating systems. All Power operating system environments that are using SAS adapters with write cache require the cache to be protected by using pairs of adapters.
- It is recommended for SAS cabling ease that the EXP24SX drawer be located in the same rack in which the PCIe adapter is located. However, it is often a good availability practice to split a SAS adapter pair across two PCIe drawers/nodes for availability and that may make the SAS cabling ease recommendation difficult or impossible to implement.
- HDDs and SSDs that were previously located in POWER8^(R) system units or in feature 5802 or 5803 12X-attached I/O drawers (SFF-1 bays) can be retrayed and placed in EXP24S drawers. See feature conversions previously announced on the POWER8 servers. Ordering a conversion ships an SFF-2 tray or carriage onto which you can place your existing drive after removing it from the existing SFF-1 tray/carriage. The order also changes the feature number so that IBM configuration tools can better interpret what is required.
- A maximum of 16 EXP24SX drawers is needed per PCIe drawer (#EMX0) to enable SAS cables to be properly handled by the feature EMX0 cable management bracket.

DVD and boot devices

A device capable of reading a DVD may be attached to the system and available to perform operating system installation, maintenance, problem determination, and service actions such as maintaining system firmware and I/O microcode at their latest levels. In addition, the system must be attached to a network with

software such as AIX NIM server or Linux Install Manager configured to perform these functions:

1. Disk or SSD located in an EXP24S drawer attached to a PCIe adapter
 2. A network through LAN adapters
 3. A SAN attached to a Fibre Channel (FC) or FC over Ethernet adapters and indicated to the server by the 0837 specify feature
- Assuming option 1 above, the minimum system configuration requires at least one SAS disk drive in the system for AIX and Linux and two for IBM i. If you are using option 3 above, a disk or SSD drive is not required.
 - For IBM i, a DVD drive must be available on the server when required.
 - A DVD can optionally be in the system control unit, or one or more DVDs can be located in an external enclosure such as a 7226-1U3 multimedia drawer.

Racks

The Power E1080 server is designed to fit a standard 19-inch rack. IBM Development has tested and certified the system in the IBM Enterprise rack (7965-S42). You can choose to place the server in other racks if you are confident those racks have the strength, rigidity, depth, and hole pattern characteristics required. You should work with IBM Service to determine the appropriateness of other racks.

It is highly recommended that the Power E1080 server be ordered with an IBM 42U enterprise rack (7965-S42). An initial system order is placed in a 7965-S42 rack. This is done to ease and speed client installation, provide a more complete and higher quality environment for IBM Manufacturing system assembly and testing, and provide a more complete shipping package.

The 7965-S42 is a two-meter enterprise rack that provides 42U or 42 EIA of space. Clients who don't want this rack can remove it from the order, and IBM Manufacturing will then remove the server from the rack after testing and ship the server in separate packages without a rack. Use the factory-deracking feature ER21 on the order to do this.

Front door options supported with Power E1080 system nodes for the 42U slim enterprise rack (7965-S42), front acoustic door #ECRA, high-end appearance front door #ECRF, cost-effective plain front door #ECRM.

Recommendation: The 7965-S42 has optimized cable routing, so all 42U may be populated with equipment.

The 7965-S42 rack does not need 2U on either the top or bottom for cable egress.

The system control unit is located below system node 1, with system node 1 on top of it, system node 2 on top of that, and so on.

With the two-meter 7965-S42, a rear rack extension of 12.7 cm (5 in.) (#ECRK) provides space to hold cables on the side of the rack and keep the center area clear for cooling and service access.

Recommendation: Include the above extensions when approximately more than 16 I/O cables per side are present or may be added in the future, when using the short-length, thinner SAS cables, or when using thinner I/O cables, such as Ethernet. If you use longer-length, thicker SAS cables, fewer cables will fit within the rack.

SAS cables are most commonly found with multiple EXP24SX SAS drawers (#ESLS) driven by multiple PCIe SAS adapters. For this reason, it is good practice to keep multiple EXP24SX drawers in the same rack as the PCIe I/O drawer or in a separate rack close to the PCIe I/O drawer, using shorter, thinner SAS cables. The feature ECRK extension can be good to use even with smaller numbers of cables because it enhances the ease of cable management with the extra space it provides.

Multiple service personnel are required to manually remove or insert a system node drawer into a rack, given its dimensions and weight and content.

Recommendation: To avoid any delay in service, obtain an optional lift tool (#EB2Z). One feature EB2Z lift tool can be shared among many servers and I/O drawers. The EB2Z lift tool provides a hand crank to lift and position up to 159 kg (350 lb). The EB2Z lift tool is 1.12 meters x 0.62 meters (44 in. x 24.5 in.). Note that a single system node can weigh up to 86.2 kg (190 lb).

A lighter, lower-cost lift tool is FC EB3Z** (lift tool) and EB4Z** (angled shelf kit for lift tool). The EB3Z** lift tool provides a hand crank to lift and position a server up to 400 lb. Note that a single system node can weigh up to 86.2 kg (190 lb).

PCIe I/O expansion drawer and racks

IBM Manufacturing can factory-integrate the PCIe I/O expansion drawer (#EMX0) with new server orders. Because expansion drawers complicate the access to vertical PDUs if located at the same height, IBM recommends accommodating PDUs horizontally on racks that have one or more PCIe I/O expansion drawers. Following this recommendation, IBM Manufacturing will always assemble the integrated rack configuration with horizontally mounted PDUs unless CRSP (#0469) is on the order. When specifying CRSP, you must provide the locations where the PCIe I/O expansion drawers should be placed and avoid locating them adjacent to vertical PDU locations EIA 6 through 16 and 21 through 31.

Additional PCIe I/O drawers (#EMX0) for an already installed server can be MES ordered with or without a rack. When you want IBM Manufacturing to place these MES I/O drawers into a rack and ship them together (factory integration), then the racks should be ordered as features on the same order as the I/O drawers. Regardless of the rack-integrated system to which the PCIe I/O expansion drawer is attached to, if the expansion drawer is ordered as factory-integrated, the PDUs in the rack will be defaulted to be placed horizontally to enhance cable management. Vertical PDUs can be used only if CRSP (#0469) is on the order.

After the rack with expansion drawers is delivered, you may rearrange the PDUs from horizontal to vertical. However, the IBM configurator tools will continue to assume the PDUs are placed horizontally for the matter of calculating the free space still available in the rack for additional future orders.

Power distribution units (PDUs)

- The Power E1080 server factory that is integrated into an IBM rack uses horizontal PDUs located in the EIA drawer space of the rack instead of the typical vertical PDUs found in the side pockets of a rack. This is done to aid cable routing. Each horizontal PDU occupies 1U. Vertically mounting the PDUs to save rack space can cause cable routing challenges and interfere with optimal service access.
- When mounting the horizontal PDUs, it is a good practice to place them almost at the top or almost at the bottom of the rack, leaving 2U or more of space at the very top or very bottom open for cable management. Mounting a horizontal PDU in the middle of the rack is generally not optimal for cable management.
- Two possible PDU ratings are supported: 60A (orderable in most countries) and 30A.
 - The 60A PDU supports four system node power supplies and one I/O expansion drawer or eight I/O expansion drawers.
 - The 30A PDU supports two system node power supplies and one I/O expansion drawer or four I/O expansion drawers.
- Rack-integrated system orders require at least two of either feature 7109, 7188, or 7196.
 - Intelligent PDU with universal UTG0247 connector (#7109) is for an intelligent AC power distribution unit (PDU+) that enables users to monitor the amount of power being used by the devices that are plugged in to this PDU+. This AC power distribution unit provides 12 C13 power outlets. It receives power through a UTG0247 connector. It can be used for many different countries and

applications by varying the PDU-to-wall power cord, which must be ordered separately. Each PDU requires one PDU-to-wall power cord. Supported power cords include the following features: 6489, 6491, 6492, 6653, 6654, 6655, 6656, 6657, 6658, or 6667.

- The PDU (#7188) mounts in a 19-inch rack and provides 12 C13 power outlets. The PDU has six 16A circuit breakers, with two power outlets per circuit breaker. System units and expansion units must use a power cord with a C14 plug to connect to the feature 7188. One of the following line cords must be used to distribute power from a wall outlet to the feature 7188: feature 6489, 6491, 6492, 6653, 6654, 6655, 6656, 6657, 6658, or 6667.
- The three-phase PDU (#7196) provides six C19 power outlets and is rated up to 48A. It has a 4.3 m (14 ft) fixed power cord to attach to the power source (IEC309 60A plug (3P+G)). A separate to-the-wall power cord is not required or orderable. Use the power cord 2.8 m (9.2 ft), drawer to wall/IBM PDU, (250V/10A) (#6665) to connect devices to this PDU. These power cords are different than the ones used on the feature 7188 and 7109 PDUs. Supported countries for the feature 7196 PDU are Antigua and Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guam, Guatemala, Haiti, Honduras, Indonesia, Jamaica, Japan, Mexico, Netherlands Antilles, Nicaragua, Panama, Peru, Puerto Rico, Surinam, Taiwan, Trinidad and Tobago, United States, and Venezuela.

System node power

- Four AC power supplies provide 2 + 2 redundant power for enhanced system availability. A system node is designed to continue functioning with just two working power supplies. A failed power supply can be hot swapped but must remain in the system until the replacement power supply is available for exchange.
- Four AC power cords are used for each system node (one per power supply) and are ordered using the AC Power Channel feature (#EMXA). The channel carries power from the rear of the system node to the hot-swap power supplies located in the front of the system node where they are more accessible for service.

System control unit power

- The system control unit is powered from the system nodes. UPIC cables provide redundant power to the system control unit. In a single node system two UPIC cables are attached to system node 1. In a two-node, three-node, or four-node system, one UPIC cable attaches to system node 1 and one UPIC cable attaches to system node 2. They are ordered with feature EFCA. Only one UPIC cable is enough to power the system control unit, and the other is in place for redundancy.

Concurrent maintenance or hot-plug options

The following options are maintenance or hot-plug capable:

- EXP24S SAS storage enclosure drawer.
- Drives in the EXP24S storage enclosure drawer.
- NVMe U.2 drives.
- PCI extender cards, optical PCIe link IO expansion card.
- PCIe I/O adapters.
- PCIe I/O drawers.
- PCIe to USB conversion card.
- SMP cables.
- System node AC power supplies: Two functional power supplies must remain installed at all times while the system is operating.
- System node fans.
- System control unit fans.

- System control unit operations panel.
- Time of Day battery.
- UPIC interface card in SCU.
- UPIC power cables from system node to system control unit.

If the system boot device or system console is attached using an I/O adapter feature, that adapter may not be hot-plugged if a nonredundant topology has been implemented.

You can access hot-plug procedures in the product documentation at [IBM Documentation](#).

Active Memory Expansion

AME 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 and 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.

AME uses CPU resource to compress and 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 and decompression.

The Power E1080 includes a hardware accelerator designed to boost AME efficiency and uses less Power core resource. You have a great deal of control over AME usage. Each individual AIX partition can turn on or turn off AME. Control parameters set the amount of expansion desired in each partition to help control the amount of CPU used by the AME function. An IPL is required for the specific partition that is turning 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 model can run the planning tool. In addition, a one-time, 60-day trial of AME is available to enable more exact memory expansion and CPU measurements. You can request the trial using the [Capacity on Demand](#) web page.

AME is enabled by chargeable hardware feature EM89, which can be ordered with the initial order of the system 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 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 AME. Normal licensing requirements apply.

Active Memory Mirroring

Active Memory Mirroring (AMM) is available to enhance resilience by mirroring critical memory used by the PowerVM hypervisor, so that it can continue operating in the event of a memory failure.

IBM i operating system

For clients loading the IBM i operating system, the four-digit numeric QPRCFEAT value is generally the same as the four-digit numeric feature number for the processors used in the system. The Power E1080 3.9 GHz processor feature is an exception to this rule. For the Power E1080:

Feature	Description
EDP2	Processor (3.65 - 3.9 GHz 40-core node) - QPRCFEAT value for the system is EDP2
EDP3	Processor (3.6 - 4.15 GHz 48-core node) - QPRCFEAT value for the system is EDP3
EDP4*	Processor (3.55 - 4.0 GHz 60-core node) - QPRCFEAT value for the system is EDP4

Power E1080 is IBM i software tier P30.

If the 5250 Enterprise Enablement function is to be used on the server, order one or more feature ED2Z or order the full system 5250 enablement feature ED30. Feature ED2Z provides one processor core's worth of 5250 capacity, which can be spread across multiple physical processor cores or multiple partitions.

Capacity Backup for IBM i

The Capacity Backup (CBU) designation can help meet your requirements for a second system to use for backup, high availability, and disaster recovery. It enables you to temporarily transfer IBM i processor license entitlements and 5250 Enterprise Enablement entitlements purchased for a primary machine to a secondary CBU-designated system. Temporarily transferring these resources instead of purchasing them for your secondary system may result in significant savings. Processor activations cannot be transferred as part of this CBU offering, however programs such as Power Enterprise Pools are available for moving or sharing processor activations.

The CBU specify feature number 4891 is available only as part of a new server purchase to a 9080-HEX. Certain system prerequisites must be met, and system registration and approval are required before the CBU specify feature can be applied on a new server. A used system that has an existing CBU feature cannot be registered. The only way to attain a CBU feature that can be registered is with a plant order.

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

After a CBU system designation is approved and the system is installed, you can temporarily move your IBM i processor license entitlements and 5250 Enterprise Enablement entitlements from the primary system to the CBU system. The CBU system can then better support fail-over and role swapping for a full range of test, disaster recovery, and high availability scenarios. Temporary entitlement transfer means that the entitlement transfer from the primary system to the CBU system and may remain in use on the CBU system as long as the registered primary and CBU system are in deployment for the high availability or disaster recovery operation. The primary system for a Power E1080 server can be any of the following:

- 9080-HEX
- 9080-M9S

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

Before you can temporarily transfer IBM i processor license entitlements from the registered primary system, you must have more than one IBM i processor license on the primary machine and at least one IBM i processor license on the CBU server. An activated processor must be available on the CBU server to use the transferred entitlement. You may then transfer any IBM i processor entitlements above the minimum of one entitlement (more may be required depending on the replication technology), assuming the total IBM i workload on the primary system does not require the IBM i entitlement you would like to transfer during the time of the transfer. During this temporary transfer, you may see IBM i license "Out of

Compliance" warning messages from the CBU system. Such messages that arise in the situation of the temporarily transferred IBM i entitlements Machine may be ignored.

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

For example, if you have a 64-core Power E980 as your primary system with twenty IBM i processor license entitlements (nineteen above the minimum) and two 5250 Enterprise Enablement entitlements (one above the minimum), you can temporarily transfer up to nineteen IBM i entitlements and one 5250 Enterprise Enablement.

If your primary or CBU machine is sold or discontinued from use, any temporary entitlement transfers must be returned to the machine on which they were originally acquired. For CBU registration and further information, see the [Capacity Backup](#) website.

* Features EDP4, EDPD, and ELCM are not available to order in China.

**Features EB3Z and EB4Z are not available to order in Albania, Bahrain, Bulgaria, Croatia, Egypt, Greece, Jordan, Kuwait, Kosovo, Montenegro, Morocco, Oman, UAE, Qatar, Saudi Arabia, Serbia, Slovakia, Slovenia, Taiwan, and Ukraine.

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 E1080 system RAS

The Power E1080 comes with dual line cord redundancy along with n+1 power supply redundancy and n+1 fan rotor redundancy. Power supplies and fans are concurrently maintainable.

The system service processor is redundant with a dynamic failover capability. Each system node has dual processor clock logic with a dynamic failover capability.

Concurrent maintenance of the real-time clock battery and the operator panel is also provided.

Memory subsystem RAS

New Differential DDIMM. Contains new memory buffer. Features OMI connection to processor with retry capabilities and dual lanes with lane reduction.

Key features

- Retry operations
- Each connection has two links with half bandwidth mode
- DDIMM supports chipkill correction plus spare DRAMS, a dynamic row repair capability
- 4U DDIMM includes a spare PMIC (or redundant) PMIC (Power Management Interface controller)

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 IBM PowerVM environment, it will perform a reset/reload if it detects the loss of the service processor.

Environmental monitoring functions

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

Power10 processor functions

As in Power9™, the Power10 processor has the ability to do processor instruction retry for some transient errors and provide core-contained checkstops for certain solid faults.

Cache availability

The L3 caches in the Power10 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 if the data is not owned by the hypervisor.

PCI extended error handling

PCI extended error handling (EEH)-enabled adapters respond to a special data packet generated from the affected PCI slot hardware by calling system firmware, which will examine the affected bus, allow the device driver to reset it, and continue without a system reboot. For Linux, EEH support extends to the majority of frequently used devices, although some third-party PCI devices may not provide native EEH support.

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 SSR Setup, Install, and Service
- Error Detection and Fault Isolation (ED/FI)
- First Failure Data Capture (FFDC)
- Guiding Light service indicators

- Service labels and service diagrams available on the system and delivered through IBM Documentation
- Step-by-step service procedures documented in IBM Documentation or available through the HMC
- Automatic reporting of serviceable events to IBM through the Electronic Service Agent Call Home application

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 Power10 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.

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 guiding light LED implementation, the system can clearly identify components for replacement by using specific component-level LEDs . The servicer can 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 re-create 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.

Standalone 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 :

- EXP24S SAS Storage enclosure drawer.
- Drives in the EXP24S Storage enclosure drawer.
- NVMe U.2 drives.
- PCI extender cards, Optical PCIe Link IO Expansion Card.
- PCIe I/O adapters.
- PCIe I/O drawers.
- PCIe to USB Conversion card.
- SMP cables.
- System node AC power supplies: Two functional power supplies must remain installed at all times while the system is operating.
- System node fans.
- System control unit fans.
- System control unit Op Panel.
- Time of Day battery.
- UPIC Interface card in SCU.
- UPIC power cables from system node to system control unit.

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 IBM i 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. Client 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 to another Power 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 E1080 and other servers using Power10-technology processors with firmware level FW1010 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 Power10 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. It runs on its own boundary and does not require resources from a system processor to be operational to perform its tasks.

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 if the system i OS is running, it can also be configured to report OS detected errors. Call Home on the HMC will continue to report platform errors. 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](#).

Statement of general direction

IBM intends to deploy Power10 in select Power Virtual Server datacenters. Power10 in Power Virtual Server is intended to deliver improved performance, scale, security, and embedded AI capabilities, allowing clients to further enhance their Power hybrid cloud infrastructure.

IBM intends to offer pay-per-use metering of Red Hat OpenShift^(R) environments running as an integral part of the Power Systems Private Cloud Solution with Dynamic Capacity. This enablement is designed to expand the breadth of applications that may leverage shared resources and instantaneous access to pay-per-use capacity that the Power Private Cloud Solution can provide.

Information concerning Red Hat products are obtained from Red Hat. Questions concerning these products should be directed to Red Hat. Red Hat products are offered to users under Red Hat's subscription terms and conditions. Availability and support are the responsibility of Red Hat. IBM is not liable and makes no warranties, express or implied, regarding Red Hat products, including but not limited to any implied warranties or condition of merchantability and fitness for a particular purpose. Moreover, all statements regarding IBM or Red Hat future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only. Information regarding potential future third-party products that may work with IBM products should not be relied on in making a purchase decision. The information mentioned regarding potential future third-party products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future third-party products may not be incorporated into any contract. The development, release, and timing of any future features of functionality described from IBM or Red Hat remains at the sole discretion of IBM or Red Hat, as applicable.

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 9080 machine type:

Description	Machine type	Model number	Feature number
IBM Power E1080	9080	HEX	
QSFP28 to SFP28 Connector	9080	HEX	EB49
QSFP+ 40GbE Base-SR4 Transceiver	9080	HEX	EB57
IBM Rack-mount Drawer Bezel and Hardware	9080	HEX	EBAB
OEM Rack-mount Drawer Bezel and Hardware	9080	HEX	EBAC
E1080 Control Unit External USB Enablement Cable	9080	HEX	EC6N
800GB Mainstream NVMe U.2 SSD 4k for AIX/Linux	9080	HEX	EC7Q
4.3m (14-Ft) PDU to wall 3PH/40A 200-240V Power Cord	9080	HEX	ECJ6
Single 5250 Enterprise Enablement	9080	HEX	ED2Z
Full 5250 Enterprise Enablement	9080	HEX	ED30
100 GB DDR4 Mobile Memory Activation for HEX Feature EDAD not orderable in China	9080	HEX	EDAB
256 GB Base Memory Activation (Pools 2.0) Feature EDAG not orderable in China	9080	HEX	EDAG

512 GB Base Memory Activation (Pools 2.0) Feature EDAH not orderable in China	9080	HEX	EDAH
256 GB Base Memory Activation Linux only Feature EDAL not orderable in China	9080	HEX	EDAL
512 GB Base Memory Activation Linux only Feature EDAM not orderable in China	9080	HEX	EDAM
1 GB Base Memory activation (Pools 2.0) from Static Feature EDAP not orderable in China	9080	HEX	EDAP
100 GB Base Memory activation (Pools 2.0) from Static Feature EDAQ not orderable in China	9080	HEX	EDAQ
512 GB Base Memory activation (Pools 2.0) from Static Feature EDAR not orderable in China	9080	HEX	EDAR
500 GB Base Memory activation (Pools 2.0) from Static Feature EDAS not orderable in China	9080	HEX	EDAS
1 GB Base Memory activation (Pools 2.0) MES only Feature EDAT not orderable in China	9080	HEX	EDAT
100 GB Base Memory activation (Pools 2.0) MES only Feature EDAU not orderable in China	9080	HEX	EDAU
100 GB Base Memory Activation (Pools 2.0) from Mobile Feature EDAV not orderable in China	9080	HEX	EDAV
500 GB Base Memory Activation (Pools 2.0) from Mobile Feature EDAW not orderable in China	9080	HEX	EDAW
512 GB Base Memory Activation Linux only - Conversion Feature EDAX not orderable in China	9080	HEX	EDAX
2 x SMP cable brackets for non-IBM Rack	9080	HEX	EDBK
Flexible service processor	9080	HEX	EDFP
5U System node Indicator drawer	9080	HEX	EDN1
40-core (4x10) 3.65 to 3.90 GHz (max) Power10 Processor with 5U system node drawer	9080	HEX	EDP2
48-core (4x12) 3.60 to 4.15 GHz (max) Power10 Processor with 5U system node drawer	9080	HEX	EDP3
60-core (4x15) 3.55 to 4.00 GHz (max) Power10 Processor with 5U system node drawer Feature EDP4 not orderable in China	9080	HEX	EDP4
1 core Processor Activation for #EDP2	9080	HEX	EDPB
1 core Processor Activation for #EDP3	9080	HEX	EDPC
1 core Processor Activation for #EDP4 Feature EDPD not orderable in China	9080	HEX	EDPD
Mobile processor activation for HEX	9080	HEX	EDPZ
System Node to System Control Unit Cable Set for Drawer 2	9080	HEX	EFCE
System Node to System Control Unit Cable Set for Drawer 3	9080	HEX	EFCF
System Node to System Control Unit Cable Set for Drawer 4	9080	HEX	EFCG
System Node to System Control Unit Cable Set for Drawer 1	9080	HEX	EFCH
PCIe x16 to CXP Converter Card, Supports optical cables	9080	HEX	EJ24
4-NVMe U.2 (7mm) Flash drive bays	9080	HEX	EJBC
4.3m (14-Ft) PDU to wall 24A 200-240V Power Cord North America	9080	HEX	ELC1
4.3m (14-Ft) PDU to wall 3PH/24A 415V Power Cord North America	9080	HEX	ELC2
Power Linux processor activation for #EDP2	9080	HEX	ELCL
Power Linux processor activation for #EDP4 Feature ELCM not orderable in China	9080	HEX	ELCM
PowerVM for Linux indicator for HEX	9080	HEX	ELCN
Power Linux processor activation for #EDP3	9080	HEX	ELCQ
512 GB Power Linux Memory Activations for HEX	9080	HEX	ELME
Active Memory expansion enablement for HEX	9080	HEX	EM8F
1 GB Memory activation for HEX	9080	HEX	EMAZ
500 GB DDR4 Mobile Memory Activation for HEX	9080	HEX	EMBK
512 GB Memory Activations for HEX	9080	HEX	EMBZ
128 GB (4x32GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory	9080	HEX	EMC1

256 GB (4x64GB) DDIMMs, 3200 MHZ, 16GBIT DDR4 Memory	9080	HEX	EMC2
512 GB (4x128GB) DDIMMs, 2933 MHZ, 16GBIT DDR4 Memory	9080	HEX	EMC3
1 TB (4x256GB) DDIMMs, 2933 MHZ, 16GBIT DDR4 Memory	9080	HEX	EMC4
100 GB of #EMAZ Memory activation for HEX	9080	HEX	EMQZ
1 core Base Processor Activation (Pools 2.0) for EDP2 any OS Feature EPDC not orderable in China	9080	HEX	EPDC
1 core Base Processor Activation (Pools 2.0) for EDP3 any OS Feature EPDD not orderable in China	9080	HEX	EPDD
1 core Base Processor Activation (Pools 2.0) for EDP4 any OS Feature EPDS not orderable in China	9080	HEX	EPDS
1 core Base Processor Activation (Pools 2.0) for EDP2 Linux only Feature EPDU not orderable in China	9080	HEX	EPDU
1 core Base Processor Activation (Pools 2.0) for EDP3 Linux only Feature EPDW not orderable in China	9080	HEX	EPDW
1 core Base Processor Activation (Pools 2.0) for EDP4 Linux only Feature EPDX not orderable in China	9080	HEX	EPDX
1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Static) Feature EPS0 not orderable in China	9080	HEX	EPS0
1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Static) Feature EPS1 not orderable in China	9080	HEX	EPS1
1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Static) Feature EPS2 not orderable in China	9080	HEX	EPS2
1 core Base Proc Act (Pools 2.0) for #EDP2 Linux (from Static) Feature EPS5 not orderable in China	9080	HEX	EPS5
1 core Base Proc Act (Pools 2.0) for #EDP3 Linux (from Static) Feature EPS6 not orderable in China	9080	HEX	EPS6
1 core Base Proc Act (Pools 2.0) for #EDP4 Linux (from Static) Feature EPS7 not orderable in China	9080	HEX	EPS7
1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Mobile Prev) Feature EPSK not orderable in China	9080	HEX	EPSK
1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Mobile Prev) Feature EPSL not orderable in China	9080	HEX	EPSL
1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Mobile Prev) Feature EPSM not orderable in China	9080	HEX	EPSM
AIX Update Access Key (UAK)	9080	HEX	ESWK

The following are features already announced for the IBM Power Systems 9080 machine type:

Description	Machine type	Model number	Feature number
Solution Delivery Integration (SDI) Bulk MES Order Indicator	9080	HEX	0003
EMEA Bulk MES Indicator	9080	HEX	0004

OEM Light Manufacturing Order Indicator - EMPTY	9080	HEX	0006
One CSC Billing Unit	9080	HEX	0010
Ten CSC Billing Units	9080	HEX	0011
Mirrored System Disk Level, Specify Code	9080	HEX	0040
Device Parity Protection-All, Specify Code	9080	HEX	0041
Device Parity RAID-6 All, Specify Code	9080	HEX	0047
Special Manufacturing Operations Indicator	9080	HEX	0098
RISC-to-RISC Data Migration	9080	HEX	0205
AIX Partition Specify	9080	HEX	0265
Linux Partition Specify	9080	HEX	0266
IBM i Operating System Partition Specify	9080	HEX	0267
Specify Custom Data Protection	9080	HEX	0296
Mirrored Level System Specify Code	9080	HEX	0308
RAID Hot Spare Specify	9080	HEX	0347
HMC Factory Integration Specify	9080	HEX	0374
Display Factory Integration Specify	9080	HEX	0375
Reserve Rack Space for UPS	9080	HEX	0376
Reserve Rack Space for HMC	9080	HEX	0377
Reserve Rack Space for Display	9080	HEX	0378
SAN Load Source Specify	9080	HEX	0837
USB 500 GB Removable Disk Drive	9080	HEX	1107
Decline Electronic Service Agent Install Indicator	9080	HEX	1120
Custom Service Specify, Rochester Minn, USA	9080	HEX	1140
300GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)	9080	HEX	1953
600GB 10k RPM SAS SFF-2 HDD for AIX/Linux	9080	HEX	1964
Primary OS - IBM i	9080	HEX	2145
Primary OS - AIX	9080	HEX	2146
Primary OS - Linux	9080	HEX	2147
SAS Cable (AE) Adapter to Enclosure, single controller/single path 3M	9080	HEX	3684
Rack Indicator- Not Factory Integrated	9080	HEX	4650
Rack Indicator, Rack #1	9080	HEX	4651
Rack Indicator, Rack #2	9080	HEX	4652
Rack Indicator, Rack #3	9080	HEX	4653
Rack Indicator, Rack #4	9080	HEX	4654
Rack Indicator, Rack #5	9080	HEX	4655
Rack Indicator, Rack #6	9080	HEX	4656
Rack Indicator, Rack #7	9080	HEX	4657
Rack Indicator, Rack #8	9080	HEX	4658
Rack Indicator, Rack #9	9080	HEX	4659
Rack Indicator, Rack #10	9080	HEX	4660
Rack Indicator, Rack #11	9080	HEX	4661
Rack Indicator, Rack #12	9080	HEX	4662
Rack Indicator, Rack #13	9080	HEX	4663
Rack Indicator, Rack #14	9080	HEX	4664
Rack Indicator, Rack #15	9080	HEX	4665
Rack Indicator, Rack #16	9080	HEX	4666
CBU SPECIFY	9080	HEX	4891
Software Preload Required	9080	HEX	5000
PowerVM Enterprise Edition	9080	HEX	5228
PCIe2 LP 4-port 1GbE Adapter	9080	HEX	5260
Sys Console On HMC	9080	HEX	5550
PCIe2 4-port 1GbE Adapter	9080	HEX	5899
Power Cord 4.3m (14-ft), Drawer to IBM PDU (250V/10A)	9080	HEX	6458
Power Cord 4.3m (14-ft), Drawer To OEM PDU (125V, 15A)	9080	HEX	6460
Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (250V/15A) U. S.	9080	HEX	6469
Power Cord 1.8m (6-ft), Drawer to Wall (125V/15A)	9080	HEX	6470
Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU	9080	HEX	6470

(250V/10A)	9080	HEX	6471
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/16A)	9080	HEX	6472
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/10A)	9080	HEX	6473
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/13A)	9080	HEX	6474
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/16A)	9080	HEX	6475
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9080	HEX	6476
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/16A)	9080	HEX	6477
Power Cord 2.7 M(9-foot), To wall/OEM PDU, (250V, 16A)	9080	HEX	6478
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (125V/15A or 250V/10A)	9080	HEX	6488
4.3m (14-Ft) 3PH/32A 380-415V Power Cord	9080	HEX	6489
4.3m (14-Ft) 1PH/63A 200-240V Power Cord	9080	HEX	6491
4.3m (14-Ft) 1PH/60A (48A derated) 200-240V Power Cord	9080	HEX	6492
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9080	HEX	6493
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9080	HEX	6494
Power Cord 2.7M (9-foot), To wall/OEM PDU, (250V, 10A)	9080	HEX	6496
Power Cable - Drawer to IBM PDU, 200-240V/10A	9080	HEX	6577
Power Cord 2.7M (9-foot), To wall/OEM PDU, (125V, 15A)	9080	HEX	6651
4.3m (14-Ft) 3PH/16A 380-415V Power Cord	9080	HEX	6653
4.3m (14-Ft) 1PH/30A (24A derated) Power Cord	9080	HEX	6654
4.3m (14-Ft) 1PH/30A (24A derated) WR Power Cord	9080	HEX	6655
4.3m (14-Ft) 1PH/32A Power Cord	9080	HEX	6656
4.3m (14-Ft) 1PH/32A Power Cord-Australia	9080	HEX	6657
4.3m (14-Ft) 1PH/30A (24A derated) Power Cord-Korea	9080	HEX	6658
Power Cord 2.7M (9-foot), To wall/OEM PDU, (250V, 15A)	9080	HEX	6659
Power Cord 4.3m (14-ft), Drawer to wall/OEM PDU (125V/15A)	9080	HEX	6660
Power Cord 2.8m (9.2-ft), Drawer to IBM PDU, (250V/10A)	9080	HEX	6665
4.3m (14-Ft) 3PH/32A 380-415V Power Cord-Australia	9080	HEX	6667
Power Cord 4.3M (14-foot), Drawer to OEM PDU, (250V, 15A)	9080	HEX	6669
Power Cord 2.7M (9-foot), Drawer to IBM PDU, 250V/10A	9080	HEX	6671
Power Cord 2M (6.5-foot), Drawer to IBM PDU, 250V/10A	9080	HEX	6672
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9080	HEX	6680
Intelligent PDU+, 1 EIA Unit, Universal UTG0247 Connector	9080	HEX	7109
Environmental Monitoring Probe	9080	HEX	7118
Power Distribution Unit	9080	HEX	7188
Power Distribution Unit (US) - 1 EIA Unit, Universal, Fixed Power Cord	9080	HEX	7196
OEM (Generic) Indicator	9080	HEX	7770
OEM (GROUPE BULL) Indicator	9080	HEX	7773
OEM (Hitachi) Indicator	9080	HEX	7775
OEM Publications for IBM Logo Product	9080	HEX	7779
Indicator for Inspur Power servers	9080	HEX	777A
Ethernet Cable, 15m, Hardware Management Console to System Unit	9080	HEX	7802
Base Customer Spec Plcmnt	9080	HEX	8453
Order Routing Indicator- System Plant	9080	HEX	9169
Language Group Specify - US English	9080	HEX	9300
New AIX License Core Counter	9080	HEX	9440
New IBM i License Core Counter	9080	HEX	9441
New Red Hat License Core Counter	9080	HEX	9442
New SUSE License Core Counter	9080	HEX	9443
Other AIX License Core Counter	9080	HEX	9444

Other Linux License Core Counter	9080	HEX	9445
3rd Party Linux License Core Counter	9080	HEX	9446
VIOS Core Counter	9080	HEX	9447
Other IBM i License Core Counter	9080	HEX	9448
Other License Core Counter	9080	HEX	9449
Month Indicator	9080	HEX	9461
Day Indicator	9080	HEX	9462
Hour Indicator	9080	HEX	9463
Minute Indicator	9080	HEX	9464
Qty Indicator	9080	HEX	9465
Countable Member Indicator	9080	HEX	9466
Language Group Specify - Dutch	9080	HEX	9700
Language Group Specify - French	9080	HEX	9703
Language Group Specify - German	9080	HEX	9704
Language Group Specify - Polish	9080	HEX	9705
Language Group Specify - Norwegian	9080	HEX	9706
Language Group Specify - Portuguese	9080	HEX	9707
Language Group Specify - Spanish	9080	HEX	9708
Language Group Specify - Italian	9080	HEX	9711
Language Group Specify - Canadian French	9080	HEX	9712
Language Group Specify - Japanese	9080	HEX	9714
Language Group Specify - Traditional Chinese (Taiwan)	9080	HEX	9715
Language Group Specify - Korean	9080	HEX	9716
Language Group Specify - Turkish	9080	HEX	9718
Language Group Specify - Hungarian	9080	HEX	9719
Language Group Specify - Slovakian	9080	HEX	9720
Language Group Specify - Russian	9080	HEX	9721
Language Group Specify - Simplified Chinese (PRC)	9080	HEX	9722
Language Group Specify - Czech	9080	HEX	9724
Language Group Specify - Romanian	9080	HEX	9725
Language Group Specify - Croatian	9080	HEX	9726
Language Group Specify - Slovenian	9080	HEX	9727
Language Group Specify - Brazilian Portuguese	9080	HEX	9728
Language Group Specify - Thai	9080	HEX	9729
Customer Install MES	9080	HEX	9742
Notify CSO After Install	9080	HEX	9743
Product Renovated by IBM Indicator	9080	HEX	9993
10m (30.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers	9080	HEX	EB2J
30m (90.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers	9080	HEX	EB2K
Lift Tool	9080	HEX	EB2Z
Mobile Enablement	9080	HEX	EB35
Lift tool based on GenieLift GL-8 (standard)	9080	HEX	EB3Z
Feature EB3Z not orderable in Albania, Bahrain, Bulgaria, Croatia, Egypt, Greece, Jordan, Kuwait, Kosovo, Montenegro, Morocco, Oman, UAE, Qatar, Saudi Arabia, Serbia, Slovakia, Slovenia, Taiwan and Ukraine			
10GbE Optical Transceiver SFP+ SR	9080	HEX	EB46
25GbE Optical Transceiver SFP28	9080	HEX	EB47
1GbE Base-T Transceiver RJ45	9080	HEX	EB48
0.5m SFP28/25GbE copper Cable	9080	HEX	EB4J
1.0m SFP28/25GbE copper Cable	9080	HEX	EB4K
2.0m SFP28/25GbE copper Cable	9080	HEX	EB4M
2.0m QSFP28/100GbE copper split Cable to SFP28 4x25GbE	9080	HEX	EB4P
Service wedge shelf tool kit for EB3Z	9080	HEX	EB4Z
Feature EB4Z not orderable in Albania, Bahrain, Bulgaria, Croatia, Egypt, Greece, Jordan, Kuwait, Kosovo, Montenegro, Morocco, Oman, UAE, Qatar, Saudi Arabia, Serbia, Slovakia, Slovenia, Taiwan and Ukraine			
100GbE Optical Transceiver QSFP28	9080	HEX	EB59
1.0M 100GbE Copper Cable QSFP28	9080	HEX	EB5K
1.5M 100GbE Copper Cable QSFP28	9080	HEX	EB5L
2.0M 100GbE Copper Cable QSFP28	9080	HEX	EB5M
3M 100GbE Optical Cable QSFP28 (AOC)	9080	HEX	EB5R
5M 100GbE Optical Cable QSFP28 (AOC)	9080	HEX	EB5S
10M 100GbE Optical Cable QSFP28 (AOC)	9080	HEX	EB5T
15M 100GbE Optical Cable QSFP28 (AOC)	9080	HEX	EB5U
20M 100GbE Optical Cable QSFP28 (AOC)	9080	HEX	EB5V
30M 100GbE optical Cable QSFP28 (AOC)	9080	HEX	EB5W
50M 100GbE optical Cable QSFP28 (AOC)	9080	HEX	EB5X
100M 100GbE Optical cable QSFP28 (AOC)	9080	HEX	EB5Y
IBM i 7.3 Indicator	9080	HEX	EB73
IBM i 7.4 Indicator	9080	HEX	EB74

1.6M USB Cable	9080	HEX	EBK4
PCIe3 LP 2-Port 10Gb NIC&ROCE SR/Cu Adapter	9080	HEX	EC2R
PCIe3 2-Port 10Gb NIC&ROCE SR/Cu Adapter	9080	HEX	EC2S
PCIe3 LP 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter	9080	HEX	EC2T
PCIe3 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter	9080	HEX	EC2U
Mainstream 800 GB SSD NVMe U.2 module	9080	HEX	EC5J
Mainstream 1.6 TB SSD NVMe U.2 module	9080	HEX	EC5K
Mainstream 3.2 TB SSD NVMe U.2 module	9080	HEX	EC5L
PCIe4 LP 2-port 100Gb ROCE EN LP adapter	9080	HEX	EC67
PCIe2 LP 2-Port USB 3.0 Adapter	9080	HEX	EC6J
SAS X Cable 3m - HD Narrow 6Gb 2-Adapters to Enclosure	9080	HEX	ECBJ
SAS X Cable 6m - HD Narrow 6Gb 2-Adapters to Enclosure	9080	HEX	ECBK
SAS X Cable 15m - HD Narrow 3Gb 2-Adapters to Enclosure	9080	HEX	ECBM
SAS YO Cable 1.5m - HD Narrow 6Gb Adapter to Enclosure	9080	HEX	ECBT
SAS YO Cable 3m - HD Narrow 6Gb Adapter to Enclosure	9080	HEX	ECBU
SAS YO Cable 6m - HD Narrow 6Gb Adapter to Enclosure	9080	HEX	ECBV
SAS YO Cable 10m - HD Narrow 6Gb Adapter to Enclosure	9080	HEX	ECBW
SAS AE1 Cable 4m - HD Narrow 6Gb Adapter to Enclosure	9080	HEX	ECBY
SAS YE1 Cable 3m - HD Narrow 6Gb Adapter to Enclosure	9080	HEX	ECBZ
2M Active Optical Cable Pair for PCIe3 Expansion Drawer	9080	HEX	ECCR
10M Active Optical Cable Pair for PCIe3 Expansion Drawer	9080	HEX	ECCY
20M Active Optical Cable Pair for PCIe3 Expansion Drawer	9080	HEX	ECCZ
3.0M SAS X12 Cable (Two Adapter to Enclosure)	9080	HEX	ECDJ
4.5M SAS X12 Active Optical Cable (Two Adapter to Enclosure)	9080	HEX	ECDK
10M SAS X12 Active Optical Cable (Two Adapter to Enclosure)	9080	HEX	ECDL
1.5M SAS Y012 Cable (Adapter to Enclosure)	9080	HEX	ECDT
3.0M SAS Y012 Cable (Adapter to Enclosure)	9080	HEX	ECDU
4.5M SAS Y012 Active Optical Cable (Adapter to Enclosure)	9080	HEX	ECDV
10M SAS Y012 Active Optical Cable (Adapter to Enclosure)	9080	HEX	ECDW
0.6M SAS AA12 Cable (Adapter to Adapter)	9080	HEX	ECE0
3.0M SAS AA12 Cable	9080	HEX	ECE3
4.5M SAS AA12 Active Optical Cable (Adapter to Adapter)	9080	HEX	ECE4
4.3m (14-Ft) PDU to wall 3PH/24A 200-240V Delta-wired Power Cord	9080	HEX	ECJ5
4.3m (14-Ft) PDU to wall 3PH/48A 200-240V Delta-wired Power Cord	9080	HEX	ECJ7
High Function 9xC19 Single-Phase or Three-Phase Wye PDU plus	9080	HEX	ECJJ
High Function 9xC19 PDU plus 3-Phase Delta	9080	HEX	ECJL
High Function 12xC13 Single-Phase or Three-Phase Wye PDU plus	9080	HEX	ECJN
High Function 12xC13 PDU plus 3-Phase Delta	9080	HEX	ECJQ
Custom Service Specify, Montpellier, France	9080	HEX	ECSF
Machine Kit of Parts (MKIT) Indicator	9080	HEX	ECSL
Custom Service Specify, Mexico	9080	HEX	ECSM
Custom Service Specify, Poughkeepsie, USA	9080	HEX	ECSP
Integrated Solution Packing	9080	HEX	ECSS
Optical wrap Plug	9080	HEX	ECW0
Captive Rack identifier	9080	HEX	EFCR
Mobile Enablement	9080	HEX	EH35
SAP HANA TRACKING FEATURE	9080	HEX	EHKV
Boot Drive / Load Source in EXP24SX Specify (in #ESLS or #ELLS)	9080	HEX	EHR2
SSD Placement Indicator - #ESLS/#ELLS	9080	HEX	EHS2
PCIe3 RAID SAS Adapter Quad-port 6Gb x8	9080	HEX	EJ0J
PCIe3 LP RAID SAS Adapter Quad-Port 6Gb x8	9080	HEX	EJ0M
PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8	9080	HEX	EJ10

PCIe3 LP SAS Tape/DVD Adapter Quad-port 6Gb x8	9080	HEX	EJ11
PCIe3 12GB Cache RAID PLUS SAS Adapter Quad-port 6Gb x8	9080	HEX	EJ14
PCIe3 Crypto Coprocessor BSC-Gen3 4767	9080	HEX	EJ33
PCIe3 Crypto Coprocessor BSC-Gen3 4769	9080	HEX	EJ37
Non-paired Indicator EJ14 PCIe SAS RAID+ Adapter	9080	HEX	EJRL
Non-paired Indicator EJ0L PCIe SAS RAID Adapter	9080	HEX	EJRU
Specify Mode-1 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP24SX #ESLS/ELLS	9080	HEX	EJW1
Specify Mode-1 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9080	HEX	EJW2
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJW3
Specify Mode-2 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJW4
Specify Mode-4 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJW5
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9080	HEX	EJW6
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9080	HEX	EJW7
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP24SX #ESLS/ELLS	9080	HEX	EJWA
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWB
Specify Mode-4 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWC
Specify Mode-4 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWD
Specify Mode-4 & (3)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWE
Specify Mode-1 & (2)EJ14 & (2)Y012 for EXP24SX #ESLS/ELLS	9080	HEX	EJWF
Specify Mode-2 & (2)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWG
Specify Mode-2 & (2)EJ14 & (1)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWH
Specify Mode-2 & (4)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWJ
Specify Mode-1 & (2)EJ0L & (2)Y012 for EXP24SX #ESLS/ELLS	9080	HEX	EJWP
Specify Mode-2 & (4)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWR
Specify Mode-2 & (2)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWS
Specify Mode-2 & (2)EJ0L & (1)X12 for EXP24SX #ESLS/ELLS	9080	HEX	EJWT
PDU Access Cord 0.38m	9080	HEX	ELC0
Power Cable - Drawer to IBM PDU (250V/10A)	9080	HEX	ELC5
ESKM Load Source Specify (931GB SSD SFF-2)	9080	HEX	ELKM
ESKR Load Source Specify (1.86TB SSD SFF-2)	9080	HEX	ELKR
ESKV Load Source Specify (3.72TB SSD SFF-2)	9080	HEX	ELKV
ESKZ Load Source Specify (7.44TB SSD SFF-2)	9080	HEX	ELKZ
#ESF2 Load Source Specify (1.1TB HDD SFF-2)	9080	HEX	ELT2
#ESFS Load Source Specify (1.7TB HDD SFF-2)	9080	HEX	ELTS
#ESEU Load Source Specify (571GB HDD SFF-2)	9080	HEX	ELTU
ESK9 Load Source Specify (387GB SSD SFF-2)	9080	HEX	ELU9
ESKD Load Source Specify (775GB SSD SFF-2)	9080	HEX	ELUD
ESKH Load Source Specify (1.55TB SSD SFF-2)	9080	HEX	ELUH
ESJK Load Source Specify (931GB SSD SFF-2)	9080	HEX	ELUK
#ESNL Load Source Specify (283GB HDD SFF-2)	9080	HEX	ELUL
ESJM Load Source Specify (1.86TB SSD SFF-2)	9080	HEX	ELUM
ESJP Load Source Specify (3.72TB SSD SFF-2)	9080	HEX	ELUP
#ESNQ Load Source Specify (571GB HDD SFF-2)	9080	HEX	ELUQ
ESJR Load Source Specify (7.44TB SSD SFF-2)	9080	HEX	ELUR
ES95 Load Source Specify (387GB SSD SFF-2)	9080	HEX	ELZ5
ESNB Load Source Specify (775GB SSD SFF-2)	9080	HEX	ELZB
ESNF Load Source Specify (1.55TB SSD SFF-2)	9080	HEX	ELZF
100GB Static to Mobile Memory Auto Conversion	9080	HEX	EME0
PCIe Gen3 I/O Expansion Drawer	9080	HEX	EMX0
AC Power Supply Conduit for PCIe3 Expansion Drawer	9080	HEX	EMXA
PCIe3 6-slot Fanout Module for PCIe3 Expansion Drawer	9080	HEX	EMXH

1m (3.3-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper	9080	HEX	EN01
3m (9.8-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper	9080	HEX	EN02
5m (16.4-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper	9080	HEX	EN03
PCIe2 4-Port (10Gb+1GbE) SR+RJ45 Adapter	9080	HEX	EN05
PCIe2 LP 4-Port (10Gb+1GbE) SR+RJ45 Adapter	9080	HEX	EN0T
PCIe2 LP 2-port 10/1GbE BaseT RJ45 Adapter	9080	HEX	EN0X
PCIe3 32Gb 2-port Fibre Channel Adapter	9080	HEX	EN1A
PCIe3 LP 32Gb 2-port Fibre Channel Adapter	9080	HEX	EN1B
PCIe3 16Gb 4-port Fibre Channel Adapter	9080	HEX	EN1C
PCIe3 LP 16Gb 4-port Fibre Channel Adapter	9080	HEX	EN1D
PCIe3 LP 16Gb 4-port Fibre Channel Adapter	9080	HEX	EN1F
PCIe3 LP 2-Port 16Gb Fibre Channel Adapter	9080	HEX	EN1H
PCIe4 LP 32Gb 2-port Optical Fibre Channel Adapter	9080	HEX	EN1K
PCIe3 16Gb 2-port Fibre Channel Adapter	9080	HEX	EN2A
PCIe3 LP 16Gb 2-port Fibre Channel Adapter	9080	HEX	EN2B
Power Enterprise Pools 2.0 Enablement Feature EP20 not orderable in China	9080	HEX	EP20
Lab Services Private Cloud Capacity Assessment	9080	HEX	EP2X
Deactivation of LPM (Live Partition Mobility)	9080	HEX	EPA0
Static to Mobile Processor Auto Conversion	9080	HEX	EPE0
Horizontal PDU Mounting Hardware	9080	HEX	EPTH
High Function 9xC19 PDU: Switched, Monitoring	9080	HEX	EPTJ
High Function 9xC19 PDU 3-Phase: Switched, Monitoring	9080	HEX	EPTL
High Function 12xC13 PDU: Switched, Monitoring	9080	HEX	EPTN
High Function 12xC13 PDU 3-Phase: Switched, Monitoring	9080	HEX	EPTQ
Qty 150 of #6577	9080	HEX	EQ77
Indicator, reserve 5 EIA rack space	9080	HEX	ER16
Specify Reserve 4 EIA Rack Space for PCIe3 Expansion Drawer	9080	HEX	ER1A
Field Integration of Rack and Server	9080	HEX	ER21
387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ES94
387GB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ES95
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9080	HEX	ESB2
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9080	HEX	ESB6
387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESBA
387GB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESBB
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESBG
775GB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESBH
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESBL
1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESBM
S&H - No Charge	9080	HEX	ESC0
S&H	9080	HEX	ESC9
Virtual Capacity Expedited Shipment	9080	HEX	ESCT
571GB 10K RPM SAS SFF- HDD 4K for IBM i	9080	HEX	ESEU
600GB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux	9080	HEX	ESEV
1.1TB 10K RPM SAS SFF-2 HDD 4K for IBM i	9080	HEX	ESF2
1.2TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux	9080	HEX	ESF3
1.7TB 10K RPM SAS SFF-2 HDD 4K for IBM i	9080	HEX	ESFS
1.8TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux	9080	HEX	ESFT
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9080	HEX	ESGV
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9080	HEX	ESGZ
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJ0
931GB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJ1
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJ2
1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJ3
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJ4
3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJ5
7.45TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJ6
7.45TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJ7
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJJ
931GB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJK
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJL
1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJM
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJN
3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJP
7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESJQ

7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESJR
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9080	HEX	ESK1
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9080	HEX	ESK3
387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESK8
387GB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESK9
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESKC
775GB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESKD
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESKG
1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESKH
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESKK
931GB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESKM
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESKP
1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESKR
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESKT
3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESKV
7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESKX
7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESKZ
Specify AC Power Supply for EXP12SX/EXP24SX Storage Enclosure	9080	HEX	ESLA
ESBB Load Source Specify (387GB SSD SFF-2)	9080	HEX	ESLB
ESBH Load Source Specify (775GB SSD SFF-2)	9080	HEX	ESLH
ESBM Load Source Specify (1.55TB SSD SFF-2)	9080	HEX	ESLM
EXP24SX SAS Storage Enclosure	9080	HEX	ESLS
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESNA
775GB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESNB
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9080	HEX	ESNE
1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i	9080	HEX	ESNF
283GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)	9080	HEX	ESNL
300GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	9080	HEX	ESNM
571GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)	9080	HEX	ESNQ
600GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	9080	HEX	ESNR
1TB Removable Disk Drive Cartridge	9080	HEX	EU01
Order Placed Indicator	9080	HEX	EU29
2TB Removable Disk Drive Cartridge (RDX)	9080	HEX	EU2T
ESJ1 Load Source Specify (931GB SSD SFF-2)	9080	HEX	EU41
ESJ3 Load Source Specify (1.86TB SSD SFF-2)	9080	HEX	EU43
ESJ5 Load Source Specify (3.72TB SSD SFF-2)	9080	HEX	EU45
ESJ7 Load Source Specify (7.45TB SSD SFF-2)	9080	HEX	EU47

RDX USB External Docking Station	9080	HEX	EUA4
----------------------------------	------	-----	------

Feature EUA4 is not orderable in the following countries: China, Taiwan, Australia, India, New Zealand, Japan, Armenia, Tajikistan, Turkmenistan, Uzbekistan, Azerbaijan, Belarus, Kazakhstan, Kyrgyzstan, Russia, South Africa, and Mexico.

Standalone USB DVD drive w/cable	9080	HEX	EUA5
BP Post-Sale Services: 1 Day	9080	HEX	SVBP
IBM Systems Lab Services Post-Sale Services: 1 Day	9080	HEX	SVCS
Other IBM Post-Sale Services: 1 Day	9080	HEX	SVNN
5000 Power to Cloud Reward points	9080	HEX	SVPC
Feature SVPC not orderable in China			

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

Planned Availability Date September 17, 2021

New Features

Description	Machine type	Model number	Feature number
Rack Content Specify for 9080-HEX -7E1A	7965	S42	ER44

Rack Content Specify for 9080-HEX -12EIA	7965	S42	ER45
Rack Content Specify for 9080-HEX -17EIA	7965	S42	ER46
Rack Content Specify for 9080-HEX -22EIA	7965	S42	ER47

Feature conversions

Feature Conversions

The existing components being replaced during a model or feature conversion become the property of IBM and must be returned.

Feature conversions are always implemented on a "quantity of one for quantity of one" basis. Multiple existing features may not be converted to a single new feature. Single existing features may not be converted to multiple new features.

The following conversions are available to customers:

Feature conversions for 9080-HEX adapters features:

From FC:	To FC:	Return parts
EJ07 - PCIe3 Optical Cable Adapter for PCIe3 Expansion Drawer	EJ24 - PCIe x16 to CXP Converter Card, Supports optical cables	
EJ19 - PCIe x16 to CXP Optical or CU converter Adapter for PCIe3 Expansion Drawer	EJ24 - PCIe x16 to CXP Converter Card, Supports optical cables	No

Feature conversions for 9080-HEX cable features:

From FC:	To FC:	Return parts
ECC6 - 2M Optical Cable Pair for PCIe3 Expansion Drawer	ECCR - 2M Active Optical Cable Pair for PCIe3 Expansion Drawer	
ECC8 - 10M Optical Cable Pair for PCIe3 Expansion Drawer	ECCY - 10M Active Optical Cable Pair for PCIe3 Expansion Drawer	
ECC9 - 20M Optical Cable Pair for PCIe3 Expansion Drawer	ECCZ - 20M Active Optical Cable Pair for PCIe3 Expansion Drawer	

Feature conversions for 9080-HEX global resource activation features:

From FC:	To FC:	Return parts
EB35 - Mobile Enablement	EP20 - Power Enterprise Pools 2.0 Enablement	No
EP2X - Lab Services Private Cloud Capacity Assessment	EP20 - Power Enterprise Pools 2.0 Enablement	No

Feature conversions for 9080-HEX memory features:

From FC:	To FC:	Return parts
EMQZ - 100 GB of #EMAZ Memory activation for HEX	EDAB - 100 GB DDR4 Mobile Memory Activation for HEX	No
EDA1 - 1 GB Memory activation from previous	EDAP - 1 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EMAZ - 1 GB Memory	EDAP - 1 GB Base Memory	No

activation for HEX	activation (Pools 2.0) from Static/ Mobile	
EDA2 - 100 GB Memory activation from previous	EDAQ - 100 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EMQZ - 100 GB of #EMAZ Memory activation for HEX	EDAQ - 100 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EMBZ - 512 GB Memory Activations for HEX	EDAR - 512 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EME7 - 512GB Memory Activation from previous	EDAR - 512 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EDAB - 100 GB DDR4 Mobile Memory Activation for HEX	EDAU - 100 GB Base Memory activation (Pools 2.0) MES only	No
EDA7 - 100 GB DDR4 Mobile Memory Activation for HEX (From previous)	EDAV - 100 GB Base Memory Activation (Pools 2.0) from Mobile	No
EMBK - 500 GB DDR4 Mobile Memory Activation for HEX	EDAW - 500 GB Base Memory Activation (Pools 2.0) from Mobile	No
ELME - 512 GB Power Linux Memory Activations for HEX	EDAX - 512 GB Base Memory Activation Linux only - Conversion	No

Feature conversions for 9080-HEX processor features:

From FC:	To FC:	Return parts
EDPB - 1 core Processor Activation for #EDP2	EDPZ - Mobile processor activation for HEX	No
EDPC - 1 core Processor Activation for #EDP3	EDPZ - Mobile processor activation for HEX	No
EDPD - 1 core Processor Activation for #EDP4	EDPZ - Mobile processor activation for HEX	No
EDPZ - Mobile processor activation for HEX	EPDC - 1 core Base Processor Activation (Pools 2.0) for EDP2 any OS	No
EDPZ - Mobile processor activation for HEX	EPDD - 1 core Base Processor Activation (Pools 2.0) for EDP3 any OS	No
EDPZ - Mobile processor activation for HEX	EPDS - 1 core Base Processor Activation (Pools 2.0) for EDP4 any OS	No
EDA4 - 1 core Processor Activation for #EDP2 from previous	EPS0 - 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Static/ Mobile)	No
EDPB - 1 core Processor Activation for #EDP2	EPS0 - 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Static/ Mobile)	No
EDAC - 1 core Processor Activation for #EDP3 from previous	EPS1 - 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Static/ Mobile)	No
EDPC - 1 core Processor Activation for #EDP3	EPS1 - 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Static/ Mobile)	No
EDAD - 1 core Processor Activation for #EDP4 from previous	EPS2 - 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Static/ Mobile)	No
EDPD - 1 core Processor Activation for #EDP4/	EPS2 - 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Static/ Mobile)	No
ELCL - Power Linux processor activation for #EDP2	EPS5 - 1 core Base Proc Act (Pools 2.0) for #EDP2 Linux (from Static/ Mobile)	No
ELCQ - Power Linux processor activation for #EDP3	EPS6 - 1 core Base Proc Act (Pools 2.0) for #EDP3 Linux (from Static/ Mobile)	No
ELCM - Power Linux	EPS7 - 1 core Base Proc Act	No

processor activation for #EDP4	(Pools 2.0) for #EDP4 Linux (from Static/ Mobile)	
EP2Y - 1 core Mobile Processor Activation	EPSK - 1 core Base Proc Act OS (from Prev)	No
EP2Y - 1 core Mobile Processor Activation	EPSL - 1 core Base Proc Act OS (from Prev)	No
EP2Y - 1 core Mobile Processor Activation	EPSM - 1 core Base Proc Act OS (from Prev)	No

Feature conversions for 9080-HEX rack related features:

From FC:	To FC:	Return parts
EMXF - PCIe3 6-Slot Fanout Module for PCIe3 Expansion Drawer	EMXH - PCIe3 6-Slot Fanout Module for PCIe3 Expansion Drawer	No
EMXG - PCIe3 6-Slot Fanout Module for PCIe3 Expansion Drawer	EMXH - PCIe3 6-Slot Fanout Module for PCIe3 Expansion Drawer	No

Publications

IBM Power 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

The following information is shipped with the 9080-HEX:

- Power Hardware Information DVD
- Important Notices
- Warranty Information
- License Agreement for Machine Code

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

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 Lab Services

IBM 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 Lab Services can 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 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 Lab Services](#) 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 [Services and Consulting](#) website.

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

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

Technical information

Specified operating environment

Physical specifications

Power E1080 model HEX

System node

- Width: 445 mm (17.51 in.)
- Depth: 866.95 mm (34.1 in.)
- Height: 217.5 mm (8.55 in.) 5 EIA units
- Weight: 81.6 kg (180 lb)

System control unit

- Width: 445.6 mm (17.54 in.)
- Depth: 779.7 mm (30.7 in.)
- Height: 86 mm (3.39 in.) 2 EIA units
- Weight: 22.7 kg (50 lbs)

PCIe Gen3 I/O expansion drawer

- Width: 482 mm (19 in.)
- Depth: 802 mm (31.6 in.)
- Height: 173 mm (6.8 in.) 4 EIA units
- Weight: 54.4 kg (120 lb)

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

Earthquake conformance

Power E1080 model HEX, PCIe Gen3 I/O Expansion Drawer, EXP24SX SAS Storage Enclosure, and 7965-S42 rack conform to the zone 4 earthquake requirements and objectives of Ericsson's GR-63-CORE with the rack loaded up to 13.2 kg (29 lb) per EIA on a raised floor with a maximum height of 1066.8 mm (42 in.), and with the 7965-S42 ruggedized rack kit installed.

It is highly recommended 7965-S42 feature #ECRR is ordered if the customer is installing the rack or hardware in an earthquake environment. Additional 7965-S42 features #ECRE, #ECRM, #ECRF, #ECRG, #ECRH, #ECRJ, #7188, and #ECJN also conform to the zone 4 earthquake requirements and objectives of Ericsson's GR-63-CORE.

If the customer is installing the hardware in an earthquake environment, they must consult with a structural engineer and/or a mechanical contractor to ensure the hardware used to secure the rack is sufficient to meet the requirements for their environment.

Operating environment

- Temperature: 5 to 40 degrees C (41 to 104 degrees F) allowable
- Relative humidity: 8 to 85 percent allowable
- Wet bulb (caloric value): 15355 kcal/hour or Btu
- Electrical power: up to 4.6 kVA per node
- Noise level: 78 dBa

Software requirements

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

- AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-03-2136, or later
- AIX Version 7.2 with the 7200-04 Technology Level and Service Pack 7200-04-05-2148 (planned availability February 11, 2022)

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

- AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-01-2038, or later
- AIX Version 7.2 with the 7200-04 Technology Level and Service Pack 7200-04-01-1939, or later
- AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-05-1937, or later

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

- IBM i 7.3 TR11, or later
- IBM i 7.4 TR5, or later

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

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

- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 15, Service Pack 3, or later
- SUSE Linux Enterprise Server 12, Service Pack 5, or later
- SUSE Linux Enterprise Server 15, Service Pack 3, or later
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8 for Power LE version 8.4, or later
- Red Hat Enterprise Linux 8 for Power LE, version 8.4, or later

If installing VIOS:

- VIOS 3.1.3.10, or later

The AIX System to Maps website should be referenced to identify the most recent information on recommended AIX levels for Power:

[https://www.ibm.com^{\(R\)}/support/pages/system-aix-maps](https://www.ibm.com^(R)/support/pages/system-aix-maps)

Java™ is supported on Power10 processor-based™ system. For the best use of the performance capabilities and the most recent improvements of Power10 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 those clients who want to run Java in Linux environments, see the [Linux Download information](#) website.

For those clients who want to run Java in IBM i environments, read the following planning statements:

- Java 6 is not a supported environment for IBM i 7.3 and is no longer supported in IBM i 7.2.
- For those clients who want to run Java on IBM i, refer to the [Java on IBM](#) website

Limitations

- The PCIe I/O Expansion Drawer (#EMX0) will be limited to a maximum of two.

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 Lab Services

For details on available services, contact your IBM representative or go to the [IT Infrastructure](#) 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, and/or integration of this product.

Terms and conditions

Volume orders

Contact your IBM representative.

Products - terms and conditions

Warranty period

Warranty and Additional Coverage options:	Coverage summary:
Warranty Period:*	1 year ⁽¹⁾
Service Level:	IBM On-Site Repair, 24x7 Same Day
Service Upgrade Options :	
IBM Power Expert Care Advanced	Available in 3 or 5 years IBM On-Site Repair, 24x7 Same Day SWMA Support Proactive Support
IBM Power Expert Care Premium	Available in 1 or 3 or 5 years IBM On-Site Repair, 24x7 Same Day SWMA Support Proactive Support IBM Global Total Microcode Support Services Media Retention Enhanced Accelerated Value Program

* One year

⁽¹⁾ Known exception: Turkey

Warranty Period: 2 Years

Service Level: IBM Onsite, 24x7 Same Business Day

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.

IBM Solid State Drive (SSD) and Non-Volatile Memory Express (NVMe) devices identified in this document may have a maximum number of write cycles. IBM SSD and NVMe 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.

Warranty service

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 and Tier 2 (optional) CRU

You may install a Tier 1 or 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 1 and Tier2 CRUs:

- The System Backplane of the System Node
- System Processor Module
- Heatsink
- Processor Voltage Regulator Module
- Miscellaneous Voltage Regulator Module
- Power APPS and TPM Card
- Memory Module
- Local Control Card
- Clock Card
- Service Processor Interface Card
- SMP10 Cable Defect Card
- PCIe Adapter Cassete
- PCIe Adapter
- PCIe4 Cable Adapter
- NVMe 7mm Drive Assembly
- NVMe U.2 7mm Drive

- NVMe U.2 15mm Drive
- System Node Fan
- Power Supply
- Internal SMP10 Cable - Left
- Internal SMP10 Cable - Right
- External SMP10 Cable
- SMP Cable Defect Card internal cable
- Service Processor Cable
- UPIC Cable
- USB Pass-through Cable
- Line Cord Conduit
- Rails
- Front Cover
- Cable Management Bracket
- The System Backplane of the System Control Unit
- Service Processor Card
- Time-of-Day Battery
- Power Interface Card
- Control Panel
- Vital Product Data
- System Control Unit fan

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.

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.

Warranty service

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

Note: International Warranty Service and Power Expert Care are mutually exclusive. If Machines have Power Expert Care, then these Machines are not eligible to have International Warranty Service.

IBM Power Expert Care

IBM Power Expert Care is designed to simplify and standardize the support approach on the IBM Power E1080 Enterprise Server. The simple service tiers and pricing facilitate straightforward support for IBM Power E1080 Enterprise server.

There are two tier service levels: Advanced and Premium. Clients can select the tier that better fits their needs. Each tier provides a group of selected hardware and software services to support IBM Power E1080 Enterprise server.

For more information about IBM Power Expert Care, see Announcement [ZS21-0018](#), dated September 8, 2021.

Maintenance service options

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. The following on-site response-time objectives are available as warranty service upgrades for your machine. Available offerings are:

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.

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.

Client 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.

Based upon availability, CRUs will be shipped for next business day delivery. 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.

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

The following parts and feature(s) have been designated as Tier 1 CRUs:

- The System Backplane of the System Node
- System Processor Module
- Heatsink
- Processor Voltage Regulator Module
- Miscellaneous Voltage Regulator Module
- Power APPS and TPM Card
- Memory Module
- Local Control Card
- Clock Card
- Service Processor Interface Card
- SMP10 Cable Defect Card
- PCIe Adapter Cassete
- PCIe Adapter
- PCIe4 Cable Adapter
- NVMe 7mm Drive Assembly
- NVMe U.2 7mm Drive
- NVMe U.2 15mm Drive

- System Node Fan
- Power Supply
- Internal SMP10 Cable - Left
- Internal SMP10 Cable - Right
- External SMP10 Cable
- SMP Cable Defect Card internal cable
- Service Processor Cable
- UPIC Cable
- USB Pass-through Cable
- Line Cord Conduit
- Rails
- Front Cover
- Cable Management Bracket
- The System Backplane of the System Control Unit
- Service Processor Card
- Time-of-Day Battery
- Power Interface Card
- Control Panel
- Vital Product Data
- System Control Unit fan

CRU and Machine Exchange Service

At IBM's discretion, you will receive CRU service or IBM will initiate shipment of a replacement machine to your location. You are responsible for its installation and verification of operation. You must pack the failed machine into the shipping container that contained the replacement machine and return the failed machine to IBM. Transportation charges, both ways, are paid by IBM. You may be charged for the replacement machine if IBM does not receive the failed machine within 15 days of your receipt of the replacement.

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, clients 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

No

Machine installation

Installation is performed by IBM. IBM will install the machine in accordance with the IBM installation procedures for the machine. In the United States, contact IBM at 1-800-IBM-SERV (426-7378). In other countries, contact the local IBM office.

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.

If you purchase IBM Power Expert Care Premium, IBM will implement updates of the Microcode levels on your contracted eligible IBM Power machines

For more information about IBM Power Expert Care, see Announcement [ZS21-0018](#), dated September 8, 2021.

Prices

For all local charges, contact your local IBM representative or IBM Business Partner.

IBM Global Financing

IBM Global Financing offers competitive financing to credit-qualified clients 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 client 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 clients. Rates are based on a client'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.

Financing solutions from IBM Global Financing can help you stretch your budget and affordably acquire the new product. But beyond the initial acquisition, our end-to-end approach to IT management can also help keep your technologies current, reduce costs, minimize risk, and preserve your ability to make flexible equipment decisions throughout the entire technology lifecycle.

Announcement countries

All European, Middle Eastern, and African countries, except Armenia, Belarus, Islamic Republic of Iran, Kazakhstan, Kyrgyzstan, Russian Federation, South Africa, Syrian Arab Republic, and the Crimea region of Ukraine (Crimea).

Trademarks

POWER9 is a trademark of IBM Corporation in the United States, other countries, or both.

Power, IBM, Interconnect, AIX, IBM Cloud, IBM API Connect, PowerHA, PowerVM, POWER8, Global Technology Services and ibm.com 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 worldwide basis.

Red Hat, Ansible and OpenShift are registered trademarks of Red Hat Inc. in the U.S. and other countries.

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 October 6, 2021)

Modification to Description section - Reliability, Availability, and Serviceability.

(Corrected on September 28, 2021)

Modification to Description section - Reliability, Availability, and Serviceability.

(Corrected on September 14, 2021)

Feature change in Prices section.