

IBM Power S1022 technology-based server provides optimized and cost-effective performance and scale for businesses in pursuit of IT excellence

Table of contents

2 Overview	37 Publications
2 Key requirements	39 Technical information
2 Planned availability date	47 Terms and conditions
3 Description	51 Prices
28 Statement of general direction	52 Regional availability
29 Product number	52 Corrections

At a glance

IBM^(R) Power^(R) servers are already the most reliable and secure in their class. Now, the new IBM Power S1022 (9105-22A) technology-based server extends that leadership and introduces the essential scale-out hybrid cloud platform, uniquely architected to help clients securely and efficiently scale core operational and AI applications anywhere in a hybrid cloud. Clients can encrypt all data simply without management overhead or performance impact and drive insights faster with AI at the point of data. Clients can also gain workload deployment flexibility and agility with a single hybrid cloud currency while doing more work.

The Power S1022 features include:

- IBM Power10 processors with up to 12, 24, 32, or 40 total cores per server
- Capacity Upgrade on Demand (CUoD) processor activation features
- IBM Private Cloud with Dynamic Capacity for Enterprise Pools 2.0 processor activation features
- In-core AI inferencing and machine learning with Matrix Math Accelerator (MMA) feature
- Up to 4.0 TB of system memory distributed across 32 DDR4 differential dual inline memory module (DDIMM) slots
- Transparent Memory Encryption with no additional management setup and no performance impact
- Active Memory Mirroring (AMM) for Hypervisor available as an option to enhance resilience by mirroring critical memory used by the PowerVM^(R) hypervisor
- Ten PCIe slots with eight PCIe Gen5 capable, all with concurrent maintenance
- Up to 8 NVMe U.2 flash bays provides up to 51.2 TB of high-speed storage
- 1+1 redundant hot-plug AC Titanium power supplies in each enclosure
- IBM PowerVM-integrated virtualization with minimum processing overhead

The Power S1022 supports:

- IBM AIX^(R), IBM i, Linux^(R), and VIOS environments
- Capacity Upgrade on Demand (CUoD) processor activation entitlement
- IBM Private Cloud with Dynamic Capacity for Enterprise Pools 2.0 processor activation entitlement
- IBM Power Expert Care service tiers
- IBM Power S1022 Solution Editions for Healthcare

- IBM Power Private Cloud Rack Solutions.

Overview

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, clients need to automate and accelerate critical operations, while ensuring 24x7 availability and staying ahead of cyberthreats. Clients need applications and data to be enterprise-grade everywhere, but without adding complexity and cost.

The Power S1022 can modernize your applications and infrastructure with a frictionless hybrid cloud experience to provide the agility you need for the unpredictability of today's business. The Power S1022 can help you:

- Run workloads where you need them with efficient scaling and consistent pay-for-use consumption across public and private clouds
- Use memory encryption at the processor level designed to support the Zero Trust security approach to hybrid cloud
- Accelerate insights from data through AI inferencing directly in core
- Consolidate workloads with scalability and performance that can reduce energy consumption

The Power S1022 server is designed to improve scale performance and security while delivering class-leading reliability. The enhanced performance and scale family of systems can help deliver business agility by extending mission-critical workloads across a hybrid cloud with increased flexibility.

- **Respond faster to business demands:** The Power10 processor delivers new levels of performance as compared to IBM Power9 for the same workloads without increasing energy or carbon footprint, enabling more efficient scaling. Power Private Cloud with Dynamic Capacity includes metering of IBM i, Linux, Red Hat^(R) OpenShift^(R) Container Platform, and AIX environments for flexible consumption consistently across public, private, and hybrid clouds when combined with the Power S1022.
- **Protect data from core to cloud:** The Power10 processor provides end-to-end security with a transparent memory encryption at the processor level--without management overhead or performance impact. Power10 can also help you to stay ahead of future threats with support for post-quantum cryptography and fully homomorphic encryption.
- **Streamline insights and automation:** Power10 leverages the enhanced in-core AI inferencing capability in every server with no additional specialized hardware required. You can extract insights from your most sensitive data where it resides, eliminating the time and risk of data movement.
- **Maximize availability and reliability:** Power10 ensures your business stays up and running with inherent advanced recovery and self-healing features for infrastructure redundancy and disaster recovery in IBM Cloud^(R).

Power servers are delivering results for clients all over the globe, from new digital services for banks and real-time decision-making in manufacturing to operational efficiency for engineering and electronics. See how Power servers are contributing to client success in [IBM Case Studies](#).

Key requirements

An IBM i, IBM AIX, Linux, or VIOS operating system is required. See the [Software requirements](#) section for details.

Planned availability date

- July 22, 2022, except for feature EM6Y and ELG2
- September 9, 2022, for feature ELG2
- November 18, 2022, for feature EM6Y

Availability within a country is subject to local legal requirements.

Description

The Power S1022 (9105-22A) server is a high-performance, flexible two-socket, 2U system that provides massive scalability and flexibility. It delivers extreme density in an energy-efficient design with superior reliability and resiliency. The Power S1022 server brings a secure environment that balances mission-critical traditional workloads and modernization applications to deliver a frictionless hybrid cloud experience.

Power S1022 feature summary

- Up to two dual-chip processor modules per system server:
 - 2.90--4.0 GHz, 12-core Power10 processor (#EPG9).
- Two dual-chip processor modules per system server:
 - 2.75--4.0 GHz, 16-core Power10 processor (#EPG8).
 - 2.45--3.90 GHz, 20-core Power10 processor (#EPGA).
- MMA feature helps to perform in-core AI inferencing and machine learning where data resides.
- Processor core activation features for Pools 2.0 available on a per-core basis:
 - 1 core Base Processor Activation Pools 2.0 for #EPG9 - any OS (#EUCB).
 - 1 core Base Processor Activation Pools 2.0 for #EPG8 - any OS (#EUCA).
 - 1 core Base Processor Activation Pools 2.0 for #EPGA - any OS (#EUCC).
- CUoD Static core activation features available on a per-core basis:
 - One CUoD Static Processor Core Activation for #EPG9 (#EPF9).
 - One CUoD Static Processor Core Activation for #EPG8 (#EPF8).
 - One CUoD Static Processor Core Activation for #EPGA (#EPFA).
- Up to 4.0 TB of system memory distributed across 32 DDIMM slots per system server.
- DDR4 DDIMM memory cards. DDIMMs are extremely high-performance, high-reliability, intelligent, and dynamic random access memory (DRAM) devices:
 - 32 GB (2 x 16 GB), (#EM6N).
 - 64 GB (2 x 32 GB), (#EM6W).
 - 128 GB (2 x 64 GB), (#EM6X).
 - 256 GB (2x128GB), (#EM6Y).
- AMM for Hypervisor is available as an option to enhance resilience by mirroring critical memory used by the PowerVM hypervisor.
- PCIe slots with two processors:
 - Four x16 Gen4 or x8 Gen5 half-height, half-length slots.
 - Four x8 Gen5 half-height, half-length slots (with x16 connectors).
 - Two x8 Gen4 half-height, half-length slots (with x16 connectors).
 - All PCIe slots are concurrently maintainable.
- Integrated:
 - System management using an Enterprise Baseboard Management Controller (eBMC).
 - EnergyScale technology.

- Redundant hot-swap cooling.
- Redundant hot-swap AC Titanium power supplies.
- Up to two HMC 1 GbE RJ45 ports.
- One rear USB 3.0 port.
- One front USB 3.0 port.
- Nineteen-inch rack-mounting hardware (2U).
- Optional PCIe I/O expansion drawer with PCIe slots:
 - Up to two drawers (#EMX0).
 - Each I/O drawer holds one or two 6-slot PCIe fanout modules (#EMXH).
 - Each fanout module attaches to the system node through a PCIe copper cable adapter (#EJ24).

PowerVM

PowerVM, which delivers industrial-strength virtualization for AIX and Linux environments on Power processor-based systems, provides a virtualization-oriented performance monitor, and performance statistics are available through the HMC. These performance statistics can be used to understand the workload characteristics and to prepare for capacity planning.

Processor modules

The Power10 processor is the compute engine for the next generation of Power systems and successor to the current Power9 processor. It offers superior performance on applications such as MMA facility to accelerate computation-intensive kernels, matrix multiplication, convolution, and discrete Fourier transform. To efficiently accelerate MMA operations, the Power10 processor core implements a dense math engine (DME) microarchitecture that effectively provides an accelerator for cognitive computing, machine learning, and AI inferencing workloads.

A maximum of two Power10 processors of the same type are allowed.

- One or two 12-core, typical 2.90 to 4.0 Ghz (max) processors (#EPG9) are allowed.
- Two 16-core, typical 2.75 to 4.0 Ghz (max) processors (#EPG8) are allowed.
- Two 20-core, typical 2.45 to 3.90 GHz (max) processors (#EPGA) are allowed.

The Power S1022 offers enhanced Workload Optimized Frequency for optimum performance. This mode can dynamically optimize the processor frequency at any given time based on CPU utilization and operating environmental conditions. For a description of this feature and other power management options available for this server, see the [IBM EnergyScale for Power10 Processor-Based Systems](#) website.

The following defines the allowed quantities of processor activation entitlements:

Base Processor Core Activations for Pools 2.0 (#EP20)

- From one to a maximum of twelve Base Processor Activations (Pools 2.0) for #EPG9 - any OS (#EUCB) with one processor module are allowed.
- From one to a maximum of twenty-four Base Processor Activations (Pools 2.0) for #EPG9 - any OS (#EUCB) with two processor modules are allowed.
- From one to a maximum of thirty-two Base Processor Activations (Pools 2.0) for #EPG8 - any OS (#EUCA) with two processor modules are allowed.
- From one to a maximum of forty Base Processor Activations (Pools 2.0) for #EPGA - any OS (#EUCC) with two processor modules are allowed.

Note: Base Processor Activation for Pools 2.0 features EUCB, EUCA, and EUCC are not available to order in China.

Shared Utility Capacity on Power S1022 systems provides enhanced multisystem resource sharing and by-the-minute tracking and consumption of computing resources across a collection of systems within a Power Enterprise Pools (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 processors and software.

Clients with existing Power Enterprise Pools of Power S922 or Power S924 systems can simply add Power S1022 systems into their pool and migrate to them at the rate and pace of their choosing, as any Power S922 or Power S924 and Power S1022 server may seamlessly interoperate and share compute resources within the same pool.

A Power Private Cloud Solution infrastructure consolidated onto Power S1022 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 a CMC, which integrates with local HMCs 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 processors on all systems in a pool are activated and available for use.

Base Capacity for processor resources is purchased on each Power S922, Power S924, or Power S1022 system and is then aggregated across a defined pool of systems for consumption monitoring.

Capacity Upgrade on Demand Static Processor Core Activations

- From six to a maximum of twelve CUoD Static Processor Core Activations for #EPG9 - any OS (#EPF9) with one processor module are allowed.
- From twelve to a maximum of twenty-four CUoD Static Processor Core Activations for #EPG9 - any OS (#EPF9) with two processor modules are allowed.
- From sixteen to a maximum of thirty-two CUoD Static Processor Core Activations for #EPG8 - any OS (#EPF8) with two processor modules are allowed.
- From twenty to a maximum of forty CUoD Static Processor Core Activations for #EPGA - any OS (#EPFA) with two processor modules are allowed.

Note: At least 50 percent of the total processor cores in the Power S1022 system must be static.

Conversions CUoD Static to Base Processor core for Pools 2.0

A variety of activations fit different usage and pricing options. Static activations are permanent and support any type of application environment on this server. Base processor activations are ordered against a specific server, but they can be moved to any server within the Power Pool and can support any type of application. The following defines the allowed conversions from static to base processor and activation entitlements:

From FC:	To FC:
EPF9 - One CUoD Static Processor Core Activation for #EPG9	EUCH - 1 core Base Processor Activation (Pools 2.0) for #EPG9 - any OS (Conv from EPF9)
EPF8 - One CUoD Static Processor Core Activation for #EPG8	EUCG - 1 core Base Processor Activation (Pools 2.0) for #EPG8 - any OS (Conv from EPF8)
EPFA - One CUoD Static Processor Core Activation for #EPGA	EUCJ - 1 core Base Processor Activation (Pools 2.0) for #EPGA - any OS (Conv from EPFA)

Note: Pools 2.0 feature EP20 is required.

MMA

The Power10 processor core inherits the modular architecture of the IBM Power9 processor core, but the redesigned and enhanced microarchitecture significantly increases the processor core performance and processing efficiency. The peak computational throughput is markedly improved by new execution capabilities and optimized cache bandwidth characteristics. Extra matrix math acceleration engine can deliver significant performance gains for machine learning, particularly for AI inferencing workloads.

Memory

The Power S1022 server uses next-generation DIMMs, which are high-performance, high-reliability, high-function memory cards that contain a buffer chip, intelligence, and 2666 MHz or 3200 MHz DRAM memory. DDIMMs are placed in DDIMM slots in the server system.

- A minimum 32 GB of memory is required with one processor module. All Memory DIMMs must be ordered in pairs.
- A minimum 64 GB of memory is required with two processor modules. All Memory DIMMs must be ordered in quads.
- Each DIMM feature code delivers two physical Memory DIMMs.

Plans for future memory upgrades should be taken into account when deciding which memory feature size to use at the time of initial system order.

To assist with the plugging rules, two DDIMMs are ordered using one memory feature number. Select from:

- 32 GB (2 x 16 GB) DDIMMs, 3200 MHz, 8 Gb DDR4 Memory (#EM6N)
- 64 GB (2 x 32 GB) DDIMMs, 3200 MHz, 8 Gb DDR4 Memory (#EM6W)
- 128 GB (2 x 64 GB) DDIMMs, 3200 MHz, 16 Gb DDR4 Memory (#EM6X)
- 256 GB (2 x 128 GB) DDIMMs, 2666 MHz, 16 Gb DDR4 Memory (#EM6Y)

AMM

AMM for Hypervisor is available as an option (#EM8G) 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. A portion of available memory can be proactively partitioned such that a duplicate set may be utilized upon non-correctable memory errors. This can be implemented at the granularity of DIMMs or logical memory blocks.

Power S1022 Capacity Backup (CBU) for IBM i

The Power S1022 CBU designation enables you to temporarily transfer IBM i processor license entitlements and IBM i user license entitlements purchased for a primary machine to a secondary CBU-designated system for high availability (HA) and disaster recovery (DR) operations. Temporarily transferring these resources instead of purchasing them for your secondary system may result in significant savings. Processor activations cannot be transferred.

The CBU specify feature 0444 is available only as part of a new server purchase. Certain system prerequisites must be met, and system registration and approval are required before the CBU specify feature can be applied on a new server. Standard IBM i terms and conditions do not allow either IBM i processor license entitlements or IBM i user license entitlements to be transferred permanently or temporarily. These entitlements remain with the machine they were ordered for. When you register the association between your primary and on-order CBU system, you must agree to certain terms and conditions regarding the temporary transfer.

After a new CBU system is registered as a pair with the proposed primary system and the configuration is approved, you can temporarily move your optional IBM i processor license entitlement and IBM i user license entitlements from the primary system to the CBU system when the primary system is down or while the primary system processors are inactive. The CBU system can then support failover and role

swapping for a full range of test, DR, and HA scenarios. Temporary entitlement transfer means that the entitlement is a property transferred from the primary system to the CBU system and may remain in use on the CBU system as long as the registered primary and CBU system are in deployment for the high availability or disaster recovery operation. The intent of the CBU offering is to enable regular role-swap operations.

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. To be in compliance, the CBU will be configured in a such a manner that there will be no out-of-compliance messages prior to a failover. An activated processor must be available on the CBU server to use the transferred entitlement. You can then transfer any IBM i processor entitlements above the minimum one, 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, the CBU system's internal records of its total number of IBM i processor license entitlements are not updated, and you may see IBM i license noncompliance warning messages from the CBU system. These warning messages in this situation do not mean you are not in compliance.

The minimum number of permanent entitlements on the CBU is one; however, you are required to license all permanent workload, such as replication workload. If, for example, the replication workload consumes four processor cores at peak workload, then you are required to permanently license four cores on the CBU.

The servers with P20 or higher software tiers do not have user entitlements that can be transferred, and only processor license entitlements can be transferred.

For a Power S1022 CBU which is in the P10 software tier, the following are eligible primary systems:

- Power S1024 (9105-42A) with 48, 32, 24, or 12 cores
- Power S1022 (9105-22A) with 40, 32, 24, or 12 cores
- Power S1022s (9105-22B) with 16 or 8 cores
- Power S1014 (9105-41B) with 8 cores
- Power S924 (9009-42G)
- Power S924 (9009-42A)
- Power S922 (9009-22A)
- Power S922 (9009-22G) with minimum of 8 cores
- Power S914 (9009-41A) with minimum of 6 cores
- Power S914 (9009-41G) with minimum of 6 cores

Power S1022 software (SW) tiers for IBM i

- The 12- and 24-core processor servers (#EPG9, QPRCFEAT EPG9) are IBM i SW tier P10.
- The 32-core processor server (#EPG8, QPRCFEAT EPG8) is IBM i SW tier P10.
- The 40-core processor server (#EPGA, QPRCFEAT EPGA) is IBM i SW tier P10.

During the temporary transfer, the CBU system's internal records of its total number of IBM i processor entitlements are not updated, and you may see IBM i license noncompliance warning messages from the CBU system. Prior to a temporary transfer, the CBU will be configured in such a manner that there are no out of compliance warning messages.

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, terms and conditions, and further information, see the [IBM Power Systems: Capacity BackUp](#) website.

Power S1022 and IBM i

IBM i support is provided at a price-attractive P10 software tier even though the Power S1022 has two sockets. There are limitations to the maximum size of the partition, and all I/O must be virtualized through VIOS (VIOS is required and IBM i partitions must be set to "restricted I/O"). Up to four cores (real or virtual) per IBM i partition are supported. Multiple IBM i partitions can be created and run concurrently, and each individual partition can have up to four cores.

Titanium power supply

Titanium power supplies are designed to meet the latest efficiency regulations. The S1022 has two titanium power supplies supporting a rack: 1+1 2000 watt, 200--240 volt (#EB3N).

Redundant fans

Redundant fans are standard.

Power cords

Two power cords are required. The Power S1022 server supports power cord 4.3 meter (14 feet), drawer to wall/IBM PDU (250V/10A) in the base shipment group. See the feature listing for other options.

PCIe slots

The Power S1022 server has up to eight U.2 NVMe devices and up to ten PCIe hot-plug slots with concurrent maintenance, providing excellent configuration flexibility and expandability. For more information about PCIe slots, see the rack-integrated system with I/O expansion drawer section below.

With two Power10 processor dual-chip modules (DCM), ten PCIe slots are available:

- Four x16 Gen4 or x8 Gen5 half-height, half-length slots
- Four x8 Gen5 half-height, half-length slots (with x16 connectors)
- Two x8 Gen4 half-height, half-length slots (with x16 connectors)

With one Power10 processor DCMs, five PCIe slots are available:

- One PCIe x16 Gen4 or x8 Gen5, half-height, half-length slot
- Three PCIe x8 Gen5, half-height, half-length slots (with x16 connector)
- One PCIe x8 Gen4, half-height, half-length slot (with x16 connector)

The x16 slots can provide up to twice the bandwidth of x8 slots because they offer twice as many PCIe lanes. PCIe Gen5 slots can support up to twice the bandwidth of a PCIe Gen4 slot, and PCIe Gen4 slots can support up to twice the bandwidth of a PCIe Gen3 slot, assuming an equivalent number of PCIe lanes.

At least one PCIe Ethernet adapter is required on the server by IBM to ensure proper manufacture, test, and support of the server. One of the x8 PCIe slots is used for this required adapter.

These servers are smarter about energy efficiency when cooling the PCIe adapter environment. They sense which IBM PCIe adapters are installed in their PCIe slots and, if an adapter requires higher levels of cooling, they automatically speed up fans to increase airflow across the PCIe adapters. Note that faster fans increase the sound level of the server. Higher wattage PCIe adapters include the PCIe3 SAS adapters and SSD/flash PCIe adapters (#EJ10, #EJ14, and #EJ0J).

NVMe drive slots, RDX bay, and storage backplane options

NVMe SSDs, in the 15-millimeter carrier U.2 2.5-inch form factor, are used for internal storage in the Power S1022 system. The Power S1022 supports up to 8

NVMe U.2 devices when two storage backplanes with four NVMe U.2 drive slots (#EJ1X) are ordered. Both the 7-millimeter and 15-millimeter NVMe are supported in the 15-millimeter carrier.

Cable management arm

A folding arm is attached to the server's rails at the rear of the server. The server's power cords and the cables from the PCIe adapters or integrated ports run through the arm and into the rack. The arm enables the server to be pulled forward on its rails for service access to PCIe slots, memory, processors, and so on without disconnecting the cables from the server. Approximately 1 meter (3 feet) of cord or cable length is needed for the arm.

Integrated I/O ports

There are two HMC ports and two USB 3.0 ports. The two HMC ports are RJ45, supporting 1 Gb Ethernet connections. The eBMC USB 2.0 port can be used for communication to an Uninterrupted Power Supply (UPS) or code update.

Rack-integrated system with I/O expansion drawer

Regardless of the rack-integrated system to which the PCIe Gen3 I/O expansion drawer is attached, if the expansion drawer is ordered as factory integrated, the PDUs in the rack will be placed horizontally by default to enhance cable management.

Expansion drawers complicate the access to vertical PDUs if located at the same height. IBM recommends accommodating PDUs horizontally on racks containing one or more PCIe Gen3 I/O expansion drawers.

After the rack with expansion drawers is delivered, you can rearrange the PDUs from horizontal to vertical. However, the configurator will continue to consider the PDUs as being placed horizontally for the matter of calculating the free space still available in the rack.

Vertical PDUs can be used only if CSRP (#0469) is on the order. When specifying CSRP, you must provide the locations where the PCIe Gen3 I/O expansion drawers should be placed. Note that you must avoid locating the expansion drawers adjacent to vertical PDU locations EIA 6 through 16 and 21 through 31.

The I/O expansion drawer can be migrated from a Power9 to a Power10 processor-based system. Only I/O cards supported on Power10 in the I/O expansion drawer are allowed. Clients migrating the I/O expansion drawer configuration might have one or two PCIe3 6-slot fanout modules (#EMXH) installed in the rear of the I/O expansion drawer.

For a 2U server configuration with one processor module, up to one I/O expansion drawer (#EMX0) and one fanout module (#EMXH) connected to one PCIe x16 to CXP Converter Card Adapter (#EJ24) are supported. The right PCIe module bay must be populated by a filler module.

For a 2U server configuration with two processor modules, up to two I/O expansion drawers (#EMX0) and four fanout modules (#EMXH) connected to four PCIe x16 to CXP Converter Card Adapters (#EJ24) are supported.

Limitations:

- Mixing of prior PCIe3 fanout modules (#EMXF or #EMXG) with PCIe3 fanout modules (#EMXH) in the same I/O expansion drawer is not allowed.
- PCIe x16 to CXP Converter Card Adapter (#EJ24) requires one PCIe3 x16 slot in the system unit plus a pair of copper cables (one copper pair feature, such as feature ECCS).

RDX docking station

The RDX docking station accommodates RDX removable disk cartridges of any capacity. The disk is in a protective rugged cartridge enclosure that plugs into

the docking station. The docking station holds one removable rugged disk drive or cartridge at a time. The rugged removable disk cartridge and docking station performs saves, restores, and backups similar to a tape drive. This docking station can be an excellent entry capacity and performance option.

EXP24SX SAS storage enclosure

The EXP24SX is a storage expansion enclosure with 24 2.5-inch small form factor (SFF) SAS bays. It supports up to 24 hot-plug HDDs or SSDs in only 2 EIA of space in a 19-inch rack. The EXP24SX SFF bays use SFF Gen2 (SFF-2) carriers or trays.

The EXP24SX drawer feature ESLS is supported on the Power10 scale-out servers by AIX, IBM i, Linux, and VIOS.

With AIX, Linux, or VIOS, the EXP24SX can be ordered with four sets of 6 bays (mode 4), two sets of 12 bays (mode 2), or one set of 24 bays (mode 1). With IBM i, only one set of 24 bays (mode 1) is supported. It is possible to change the mode setting in the field using software commands along with a specifically documented procedure.

Important: 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.

Four mini-SAS HD ports on the EXP24SX are attached to PCIe Gen3 SAS adapters or attached to an integrated SAS controller in a Power10 scale-out server. The following PCIe3 SAS adapters support the EXP24SX:

- PCIe3 RAID SAS Adapter quad-port 6 Gb x8 (#EJ0J)
- PCIe3 12 GB Cache RAID Plus SAS Adapter quad-port 6 Gb x8 (#EJ14)
- PCIe3 LP RAID SAS Adapter quad-port 6 Gb x8 (#EJ0M)

Earlier-generation PCIe1 or PCIe2 SAS adapters are not supported with the EXP24SX.

The attachment between the EXP24SX and the PCIe3 SAS adapters or integrated SAS controllers is through SAS YO12 or X12 cables. X12 and YO12 cables are designed to support up to 12 Gb SAS. The PCIe Gen3 SAS adapters support up to 6 Gb throughput. The EXP24SX has been designed to support up to 12 Gb throughput if future SAS adapters support that capability. All ends of the YO12 and X12 cables have mini-SAS HD narrow connectors. Cable options are:

- X12 cable: 3-meter copper (#ECDJ), 4.5-meter optical (#ECDK), 10-meter optical (#ECDL)
- YO12 cables: 1.5-meter copper (#ECDT), 3-meter copper (#ECDU)
- 1M 100 GbE Optical Cable QSFP28 (AOC) (#EB5K)
- 1.5M 100 GbE Optical Cable QSFP28 (AOC) (#EB5L)
- 2M 100 GbE Optical Cable QSFP28 (AOC) (#EB5M)
- 3M 100 GbE Optical Cable QSFP28 (AOC) (#EB5R)
- 5M 100 GbE Optical Cable QSFP28 (AOC) (#EB5S)
- 10M 100 GbE Optical Cable QSFP28 (AOC) (#EB5T)
- 15M 100 GbE Optical Cable QSFP28 (AOC) (#EB5U)
- 20M 100 GbE Optical Cable QSFP28 (AOC) (#EB5V)
- 30M 100 GbE Optical Cable QSFP28 (AOC) (#EB5W)
- 50M 100 GbE Optical Cable QSFP28 (AOC) (#EB5X)

An AA12 cable interconnecting a pair of PCIe3 12 GB cache adapters (two #EJ14) is not attached to the EXP24SX. These higher-bandwidth cables could support 12 Gb throughput if future adapters support that capability. Copper feature ECE0 is 0.6

meters long, feature ECE3 is 3 meters long, and optical AA12 feature ECE4 is 4.5 meters long.

One no-charge specify code is used with each EXP24SX I/O Drawer (#ESLS) to communicate to IBM configurator tools and IBM Manufacturing which mode setting, adapter, and SAS cable are needed. With this specify code, no hardware is shipped. The physical adapters, controllers, and cables must be ordered with their own chargeable feature numbers. There are more technically supported configurations than are represented by these specify codes. IBM Manufacturing and IBM configurator tools such as e-config only understand and support EXP24SX configurations represented by these specify codes.

Specify code	Mode	Adapter/ Controller	Cable to drawer	Environment
EJW0	Mode 1	CEC SAS Ports	2 YO12 cables	AIX/IBM i/ Linux/VIOS
EJW1	Mode 1	One (unpaired) #EJ0J/#EJ0M	1 YO12 cable	AIX/IBM i/ Linux/VIOS
EJW2	Mode 1	Two (one pair) #EJ0J/#EJ0M	2 YO12 cables	AIX/IBM i/ Linux/VIOS
EJW3	Mode 2	Two (unpaired) #EJ0J/#EJ0M	2 X12 cables	AIX/Linux/ VIOS
EJW4	Mode 2	Four (two pair) #EJ0J/ #EJ0M	2 X12 cables	AIX/Linux/ VIOS
EJW5	Mode 4	Four (unpaired) #EJ0J/#EJ0M	2 X12 cables	AIX/Linux/ VIOS
EJW6	Mode 2	One (unpaired) #EJ0J/#EJ0M	2 YO12 cables	AIX/Linux/ VIOS
EJW7	Mode 2	Two (unpaired) #EJ0J/#EJ0M	2 YO12 cables	AIX/Linux/ VIOS
EJWF	Mode 1	Two (one pair) #EJ14	2 YO12 cables	AIX/IBM i/ Linux/VIOS
EJWG	Mode 2	Two (one pair) #EJ14	2 X12 cables	AIX/Linux/ VIOS
EJWJ	Mode 2	Four (two pair) #EJ14	2 X12 cables	AIX/Linux/ VIOS

All of the above EXP24SX specify codes assume a full set of adapters and cables able to run all the SAS bays configured. The following specify codes communicate to IBM Manufacturing a lower-cost partial configuration is to be configured where the ordered adapters and cables can run only a portion of the SAS bays. The future MES addition of adapters and cables can enable the remaining SAS bays for growth. The following specify codes are used:

Specify code	Mode	Adapter/ Controller	Cable to drawer	Environment
EJWA (1/2 of EJW7)	Mode 2	One (unpaired) #EJ0J/#EJ0M	1 YO12 cables	AIX/Linux/ VIOS
EJWB (1/2 of EJW4)	Mode 2	Two (one pair) #EJ0J/#EJ0M	1 X12 cable	AIX/Linux/ VIOS
EJWC (1/4 of EJW5)	Mode 4	One (unpaired) #EJ0J/#EJ0M	1 X12 cable	AIX/Linux/ VIOS
EJWD (1/2 of EJW5)	Mode 4	Two (unpaired) #EJ0J/#EJ0M	1 X12 cable	AIX/Linux/ VIOS
EJWE (3/4 of EJW5)	Mode 4	Three (unpaired) #EJ0J/#EJ0M	2 X12 cables	AIX/Linux/ VIOS

EJWH (1/2 of EJWJ)	Mode 2	Two (one pair) #EJ14	1 X12 cable	AIX/Linux/ VIOS
--------------------	--------	----------------------	-------------	-----------------

An EXP24SX drawer in mode 4 can be attached to two or four SAS controllers and provide a great deal of configuration flexibility. For example, if using unpaired feature EJ0J adapters, these EJ0J adapters could be in the same server in the same partition, same server in different partitions, or even different servers.

An EXP24SX drawer in mode 2 has similar flexibility. If the I/O drawer is in mode 2, then half of its SAS bays can be controlled by one pair of PCIe3 SAS adapters, such as a 12 GB write cache adapter pair (#EJ14), and the other half can be controlled by a different PCIe3 SAS 12 GB write cache adapter pair or by zero-write-cache PCIe3 SAS adapters.

Note that for simplicity, IBM configurator tools such as e-config assume that the SAS bays of an individual I/O drawer are controlled by one type of SAS adapter. As a client, you have more flexibility than e-config understands.

A maximum of 24 2.5-inch SSDs or 2.5-inch HDDs are supported in the EXP24SX 24 SAS bays. There can be no mixing of HDDs and SSDs in the same mode 1 drawer. HDDs and SSDs can be mixed in a mode 2 or mode 4 drawer, but they cannot be mixed within a logical split of the drawer. For example, in a mode 2 drawer with two sets of 12 bays, one set could hold SSDs and one set could hold HDDs, but you cannot mix SSDs and HDDs in the same set of 12 bays.

The indicator feature EHS2 helps IBM Manufacturing understand where SSDs are placed in a mode 2 or a mode 4 EXP24SX drawer. On one mode 2 drawer, use a quantity of one feature EHS2 to have SSDs placed in just half the bays, and use two EHS2 features to have SSDs placed in any of the bays. Similarly, on one mode 4 drawer, use a quantity of one, two, three, or four EHS2 features to indicate how many bays can have SSDs. With multiple EXP24SX orders, IBM Manufacturing will have to guess which quantity of feature ESH2 is associated with each EXP24SX. Consider using CSP (#0456) to reduce guessing.

Two-and-a-half-inch SFF SAS HDDs and SSDs are supported in the EXP24SX. All drives are mounted on Gen2 carriers or trays and thus named SFF-2 drives.

The EXP24SX drawer has many high-reliability design points:

- SAS drive bays that support hot swap
- Redundant and hot-plug-capable power and fan assemblies
- Dual line cords
- Redundant and hot-plug enclosure service modules (ESMs)
- Redundant data paths to all drives
- LED indicators on drives, bays, ESMs, and power supplies that support problem identification
- Through the SAS adapters or controllers, drives that can be protected with RAID and mirroring and hot-spare capability

Order two ESLA features for AC power supplies. The enclosure is shipped with adjustable depth rails and can accommodate 19-inch rack depths from 59.5--75 centimeters (23.4--29.5 inches). Slot filler panels are provided for empty bays when initially shipped from IBM.

PCIe Gen3 I/O drawer cabling option

A copper cabling option (#ECCS) is available for the scale-out servers. The cable option offers a much lower-cost connection between the server and the PCIe Gen3 I/O drawer fanout modules. The currently available Active Optical Cable (AOC) offers much longer length cables, providing rack placement flexibility. Plus, AOC cables are much thinner and have tighter bend radius and thus are much easier to cable in the rack.

The 3M Copper CXP Cable Pair (#ECCS) has the same performance and same reliability, availability, and serviceability (RAS) characteristics as the AOC cables.

One copper cable length of 3 meters is offered. Note that the cable management arm of the scale-out servers requires about 1 meter of cable.

Like the AOC cable pair, the copper pair is cabled in the same manner. One cable attaches to the top CXP port in the PCIe adapter in the x16 PCIe slot in the server system unit and then attaches to the top CXP port in the fanout module in the I/O drawer. Its cable pair attaches to the bottom CXP port of the same PCIe adapter and to the bottom CXP port of the same fanout module. Note that the PCIe adapter providing the CXP ports on the server was named a PCIe3 "Optical" Cable Adapter. In hindsight, this naming was unfortunate as the adapter's CXP ports are not unique to optical. But at the time, optical cables were the only connection option planned.

Copper and AOC cabling can be mixed on the same server. However, they cannot be mixed on the same PCIe Gen3 I/O drawer or mixed on the same fanout module.

Copper cables have the same operating system software prerequisites as AOC cables.

Racks

The Power10 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). The 7965-S42 is a 2-meter enterprise rack that provides 42U or 42 EIA of space. 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 Power10 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.

Recommendation: The 7965-S42 rack 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.

With the 2-meter 7965-S42 rack, a rear rack extension of 12.7 centimeters (5 inches) feature 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 (#EB3Z). A lighter, lower-cost lift tool is feature EB3Z¹ (lift tool) and EB4Z¹ (angled shelf kit for lift tool). The EB3Z lift tool provides a hand crank to lift and position a server up to 400 pounds. Note that a single system node can weigh up to 86.2 kilograms (190 pounds).

¹ Feature EB3Z and feature 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.

High-function (switched and monitored) PDUs plus

Hardware:

- IEC 62368-1 and IEC 60950 safety standard
- A new product safety approval
- No China 5,000-meter altitude or tropical restrictions
- Detachable inlet for 3-phase delta-wired PDU with 30A, 50A, and 60A wall plugs
- IBM Technology and Qualification approved components, such as anti-sulfur resistors (ASRs)
- Ethernet 10/100/1000 Mb/s

Software:

- Internet Protocol (IP) v4 and IPv6 support
- Secure Shell (SSH) protocol command line
- Ability to change passwords over a network

PDU description	208 V 3-phase delta	200 V--240V 1-phase or 3-phase wye
High-function 12xC13	#ECJQ/#ECJP	#ECJN/#ECJM
High-function 9xC19	#ECJL/#ECJK	#ECJJ/#ECJG

These PDUs can be mounted vertically in rack-side pockets or they can be mounted horizontally. If mounted horizontally, they each use one EIA (1U) of rack space. See feature EPTH for horizontal mounting hardware, which is used when IBM Manufacturing doesn't automatically factory-install the PDU. Two RJ45 ports on the front of the PDU enable you to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off.

Recommendation: The PDU is shipped with a generic PDU password. IBM strongly urges you to change it upon installation.

Existing and new high-function (switched and monitored) PDUs have the same physical dimensions. New high-function (switched and monitored) PDUs can be supported in the same racks as existing PDUs. Mixing of PDUs in a rack on new orders is not allowed.

Also, all factory-integrated orders must have the same PDU line cord.

The PDU features ECJQ/ECJP and ECJL/ECJK with the Amphenol inlet connector require new PDU line cords:

- #ECJ5 - 4.3 meter (14-foot) PDU to Wall 3PH/24A 200-240V Delta-wired Power Cord
- #ECJ7 - 4.3 meter (14-foot) PDU to Wall 3PH/48A 200-240V Delta-wired Power Cord

No pigtail (like #ELC0) is available because an Amphenol male inline connector is unavailable.

The PDU features ECJJ/ECJG and ECJN/ECJM with the UTG624-7SKIT4/5 inlet connector use the existing PDU line cord features 6653, 6667, 6489, 6654, 6655, 6656, 6657, 6658, 6491, or 6492.

Power S1022 Solution Edition for Healthcare

Power S1022 Solution Edition for Healthcare provides a more cost-effective solution for small-to-medium hospitals ordering a new Power S1022 server. The two available configurations are:

- S1022 Solution Edition for Healthcare for Operational Databases (ODB): Configuration supports the operation core processing of the solution.
- S1022 Solution Edition for Healthcare for Enterprise Cache Protocol (ECP): Configuration supports edge computing and branch sites to enable local processing and availability, working as an extension of a central compute server.

The following features help eligible clients to order and effectively price the configuration:

- Power10 Healthcare Solution Edition indicator (#EHM1)
- 512 GB (16 x 32 GB) Memory DIMMs bundle (#EM68)

The above features are available only with a new server order. MES orders can use the regular feature numbers with their regular pricing.

Power S1022 Solution Edition for Healthcare Operational Databases (ODB) configuration is a 32-core, two-socket server running AIX. An initial hardware order includes the following components:

Feature	Description	Default	Min	Max	Rules/Comments
EHM1	Power10 Healthcare Solution Edition indicator	1	1	1	Solution indicator feature EHM1 can apply for both ODB and ECP editions and should be for Initial orders only.
EM68	512 GB memory bundle for Power10 Solution Edition	1	1	2	Reduce price bundled feature for memory is required. If additional memory is desired on an initial order, only allow maximum quantity of eight of feature (#EM6W) - 64 GB (2 x 32 GB).
EU0K	Operator panel LCD display	1	1	1	
EPG8	16-core typical 2.75 to 4.0 Ghz (max) Power10 processor	2	2	2	

Feature	Description	Default	Min	Max	Rules/ Comments
EPF8	One CUoD Static Processor Core Activation for #EPG8	32	32	32	
EC7T	800 GB mainstream NVMe U.2 SSD 4k for AIX/Linux	2	0	**	Allow other available NVMe U.2 features. Maximum quantity from standalone configuration is allowed.
EC2T	PCIe3 LP 2-Port 25/10 Gb NIC & RoCE SR/Cu adapter	2	0	**	Allowed to replace or add other from available 10 GB or higher LAN cards. Maximum quantity from standalone configuration is allowed.
EB46	10 GbE optical transceiver SFP+ SR	2	0	**	Maximum quantity from standalone configuration is allowed.
EN1K	PCIe4 LP 32 Gb two-port optical FCI adapter	2	0	**	Qty 2 of 16 GB or higher FC cards are required. Allowed to replace or add from available 16 GB or higher FC cards. Maximum quantity from standalone configuration is allowed.
EJ1X	Storage backplane with four NVMe U.2 drive slots	1	0	2	
EJUS	Front IBM bezel for eight NVMe-bays backplane rack mount	1	1	1	

Feature	Description	Default	Min	Max	Rules/ Comments
EB3N	AC Titanium power supply - s000W for server (200--240 VAC)	2	2	2	
6458	Power cord 4.3m (14-ft), drawer to IBM PDU (250V/10A)				
0265	AIX partition specify	1	1	480	
2146	Primary OS -AIX	1	1	1	
4650	Rack indicator-not factory integrated	1	1	1	
5000	Software preload required	1	0	1	
5228	PowerVM Enterprise Edition	32	32	32	
9300	Language Group Specify - US English	1	1	1	
9440	New AIX license core counter	32	32	32	
0983	US TAA compliance indicator	0	0	1	
ESC0	S&H - no charge	0	0	1	Total minimum and total maximum = 1 of either feature ESC0 or ESC6; features ESC0 and ESC6 are mutually exclusive.
ESC5	Shipping and handling	1	0	1	Total minimum and total maximum = 1 of either feature ESC0 or ESC5; features ESC0 and ESC5 are

Feature	Description	Default	Min	Max	Rules/ Comments
					mutually exclusive.
ECW0	Optical wrap plug	2	0	2	
EU19	Cable ties & labels	1	0	1	

Power S1022 Solution Edition for Healthcare Enterprise Cache Protocol (ECP) configuration is a 32-core, two-socket server running AIX as primary operating system. An initial hardware order includes the following components:

Feature	Description	Default	Min	Max	Rules/ Comments
EHM1	Power10 Healthcare Solution Edition indicator	1	1	1	Solution indicator feature EHM1 can apply for both ODB and ECP editions and should be for initial orders only.
EM68	512 GB memory bundle for Power10 Solution Edition	1	1	2	Reduce price bundled feature for memory is required. If additional memory is desired on an initial order, only allow maximum quantity of eight of features (#EM6W) - 64 GB (2 x 32 GB).
EU0K	Operator panel LCD display	1	1	1	
EPG8	16-core typical 2.75 to 4.0 Ghz (max) Power10 processor	2	2	2	
EPF8	One CUoD Static Processor Core Activation for #EPG8	32	32	32	
EC7T	800 GB mainstream NVMe U.2 SSD 4k for AIX/Linux	2	0	**	Allow other available NVMe U.2 features. Maximum quantity

Feature	Description	Default	Min	Max	Rules/ Comments
					from standalone configuration is allowed.
EC2T	PCIe3 LP 2-Port 25/10Gb NIC & RoCE SR/Cu adapter	2	0	**	Allowed to replace or add other from available 10 GB or higher LAN cards. Maximum quantity from standalone configuration is allowed.
EB46	10 GbE optical transceiver SFP+ SR	2	0	**	Maximum quantity from standalone configuration is allowed.
EN1K	PCIe4 LP 32 Gb 2-port optical FCI adapter	1	0	**	Qty 1 of 16 GB or higher FC card is required. Allowed to replace or add from available 16 GB or higher FC cards. Maximum quantity from standalone configuration is allowed.
EJ1X	Storage backplane with four NVMe U.2 drive slots	1	0	2	
EJUS	Front IBM bezel for 8 NVMe-bays backplane rack mount	1	1	1	
EB3N	AC Titanium power supply - s000W for Server (200-240 VAC)	2	2	2	
6458	Power cord 4.3m (14-ft), drawer to IBM PDU (250V/10A)				

Feature	Description	Default	Min	Max	Rules/ Comments
0265	AIX Partition Specify	1	1	480	
2146	Primary OS -AIX	1	1	1	
4650	Rack indicator-not factory integrated	1	1	1	
5000	Software preload required	1	0	1	
5228	PowerVM Enterprise Edition	32	32	32	
9300	Language Group Specify - US English	1	1	1	
9440	New AIX license core counter	32	32	32	
0983	US TAA compliance indicator	0	0	1	
ESC0	S&H - no charge	0	0	1	Total minimum and total maximum = 1 of either feature ESC0 or ESC6; features ESC0 and ESC6 are mutually exclusive.
ESC5	Shipping and handling	1	0	1	Total minimum and total maximum = 1 of either feature ESC0 or ESC5; features ESC0 and ESC5 are mutually exclusive.
ECW0	Optical wrap plug	2	0	2	
EU19	Cable ties & labels	1	0	1	

- Additional hardware components can be added as desired following normal supported configuration rules. The above is the predefined configuration.
- Additional software and maintenance can be added as desired following normal supported configuration rules.

To see if you are eligible to order this solution edition, see the [IBM Power Solution Editions for Healthcare](#) website. Also, the sales channel can register each server using this solution edition at this website.

Power Private Cloud Rack Solutions

The two available configurations are:

- IBM Power Private Cloud Rack Solution
- IBM Power Private Cloud Starter Solution

The following feature helps eligible clients to order and effectively price the selected configuration:

- IBM Power Private Cloud Rack Solution Indicator (#ELG2)

The Power Private Cloud Rack Solution is offered in a preconfigured setup; this is an optimized full stack for a production-level environment. It leverages the unique virtualization technologies of PowerVM to accommodate all the software stack required in just three servers. Additional nodes can be added on the initial configuration as needed.

A minimum configuration includes:

- Hardware stack
 - Three Power S1022 servers
 - One FlashSystem 5200 storage enclosure with a minimum of 9.6 TB
 - Two SAN24B-6 switches with 24 FC ports and industry-leading Gen 6 FC technology
 - Optional IBM Ethernet switch with high-performance Gigabit Ethernet Layer 2 and Layer 3 switch featuring 52 ports
 - One IBM Enterprise slim rack with 42 EIA units of vertical mounting space and 19-inch rack enclosure
- Software stack
 - RHEL 8 for Power10
 - IBM PowerVM Enterprise Edition
 - IBM PowerVC for Private Cloud 2.0
 - Red Hat OpenShift Container Platform
 - IBM Spectrum^(R) Scale Data Access Edition or IBM Spectrum Scale Data Management Edition

If you have Red Hat entitlements for OpenShift Container Platform or RHEL 8, they can be deselected from the solution edition in e-config. A proof of entitlement for each software license will be requested and is mandatory in order to authorize the manufacturing and shipping of the solution.

The Power Private Cloud Starter Solution includes:

- Hardware stack
 - At least one Power S1022 server node
 - Optional IBM FlashSystem^(R) 5200 storage enclosure
- Software stack
 - Red Hat Enterprise Linux (RHEL) 8 for Power9
 - PowerVM Enterprise Edition
 - Optional PowerVC for Private Cloud 2.0
 - Red Hat OpenShift Container Platform

If you have Red Hat entitlements for OpenShift Container Platform or RHEL 8, they can be deselected from the solution edition in e-config. A proof of entitlement for each software license will be requested and is mandatory in order to authorize the manufacturing and shipping of the solution.

Reliability, Availability, and Serviceability

Reliability, fault tolerance, and data correction

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 to help ensure the highest level of product quality.

The Power10 processor-based scale-out systems come with the following RAS characteristics:

- Power10 processor RAS
- Open Memory Interface, DDIMMs RAS
- Enterprise BMC service processor for system management and Service
- AMM for Hypervisor
- NVMe drives concurrent maintenance
- PCIe adapters concurrent maintenance
- Redundant and hot-plug cooling
- Redundant and hot-plug power
- Light path enclosure and FRU LEDs
- Service and FRU labels
- Client or IBM install
- Proactive support and service -- call home
- Client or IBM service

Service processor

Power10 scale-out 2S-4S systems come with a redesigned service processor based on a Baseboard Management Controller (BMC) design with firmware that is accessible through open-source industry standard APIs, such as Redfish. An upgraded ASMI web browser user interface preserves the required RAS functions while allowing the user to perform tasks in a more intuitive way.

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 power boundary and does not require resources from a system processor to be operational to perform its tasks.

The service processor supports surveillance of the connection to the HMC and to the system firmware (hypervisor). It also provides several remote power control options, environmental monitoring, reset, restart, remote maintenance, and diagnostic functions, including console mirroring. The BMC service processors menus (ASMI) can be accessed concurrently during system operation, allowing nondisruptive abilities to change system default parameters, view and download error logs, check system health.

Redfish, an industry standard for server management, enables the Power servers to be managed individually or in a large data center. Standard functions such as inventory, event logs, sensors, dumps, and certificate management are all supported with Redfish. In addition, new user management features support multiple users and privileges on the BMC via Redfish or ASMI. User management via LDAP is also supported. The Redfish events service provides a means for notification of specific critical events such that actions can be taken to correct issues. The Redfish

telemetry service provides access to a wide variety of data (eg. power consumption, ambient, core, DIMM and I/O temperatures, etc) that can be streamed on periodic intervals.

Mutual surveillance

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

Environmental monitoring functions

The Power family does ambient and over temperature monitoring and reporting. It also adjusts fan speeds automatically based on those temperatures.

Memory subsystem RAS :

The Power10 scale-out system introduces a new 2U tall DDIMM, which has new open CAPI memory interface known as OMI for resilient and fast communication to the processor. This new memory subsystem design delivers solid RAS as described below.

Power10 processor functions

As in Power9, the Power10 processor has the ability to do processor instruction retry for some transient errors and core-contained checkstop for certain solid faults. The fabric bus design with CRC and retry persists in Power10 where a CRC code is used for checking data on the bus and has an ability to retry a faulty operation.

Cache availability

The L2/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 cache from immediately causing the system to terminate. Rather, the system tags the data and determines whether it will ever be used again. If the error is irrelevant, it will not force a 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)
- Light path service indicators
- Service and FRU labels available on the system
- 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 systems 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.
- No HMC. There are two service strategies for non-HMC systems.
 - Full-system partition with PowerVM: A single partition owns all the server resources and only one operating system may be installed. The primary service interface is through the operating system and the service processor.
 - Partitioned system with NovaLink: In this configuration, the system can have more than one partition and can be running more than one operating system. The primary service interface is through the service processor.

Service interface

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

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

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

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

As the name implies, standalone 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 and NVMe drives are supported with PowerVM. Power supplies, fans, and op panel LCD are hot pluggable.

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, processors, fans, adapter cards, 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 for Power10 processor-based systems which contains each MTM service functions of interest while physically located at the server. These include things such as installation and repair instructions, reference code look up, and so on.

Packaging for service

The following service features 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, thumbscrews, 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 filesystem. When an HMC is attached in the PowerVM environment, an ELA routine analyzes the error, forwards the event to the Service Focal Point (SFP) application running on the HMC, and notifies the system administrator that it has isolated a likely cause of the system problem. The service processor event log also records unrecoverable checkstop conditions, forwards them to the SFP application, and notifies the system administrator.

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

In the HMC-managed environment, a call home service request will be initiated from the HMC and the pertinent failure data with service parts information and part locations will be sent to an IBM service organization. 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 partition running on one Power partition system to another Power system without disrupting services. The migration transfers the entire system environment, including processor state, memory, attached virtual devices, and connected users. It provides continuous operating system and application availability during planned partition outages for repair of hardware and firmware faults. The Power10 systems using Power10-technology support 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.

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 the Electronic Service Agent (ESA) imbedded in the HMC or through a version of ESA imbedded in the operating systems for non-HMC-managed or A version of ESA that runs as a standalone call home application. 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. See the next section for specific details on this application.

IBM Electronic Services

Electronic Service Agent and Client Support Portal (CSP) 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 and collects hardware and software 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. In addition, Call Home Cloud Connect Web and Mobile capability extends the common solution and offers IBM Systems related support information applicable to Servers and Storage.

Details are available here: <https://clientvantage.ibm.com/channel/ibm-call-home-connect>.

System configuration and inventory information collected by Electronic Service Agent also can be 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 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](#).

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 through the internet (HTTPS) to provide clients a single point of exit from their site. Initiation of 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.

Client Support Portal

Client Support 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.

This web portal 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 IBM ID.

For more information on how to utilize client support portal, visit the following website or contact an [IBM Systems Services Representative](#).

Statement of general direction

IBM plans to announce IBM Power10 technology-based systems for IBM Power Private Cloud Rack solution for Oracle Database (OracleDB) to deliver preintegrated configurations in Oracle Database environments in fourth quarter of 2022.

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.

Reference information

For additional information about IBM Power Expert Care extends services and support options, see announcement [LS22-0005](#), dated July 12, 2022.

For more information on Power10 scale-out servers, see Hardware Announcements: [LG22-0029](#), dated July 12, 2022; [LG22-0031](#), dated July 12, 2022; [LG22-0032](#), dated July 12, 2022; [LG22-0033](#), dated July 12, 2022; [LG22-0034](#), dated July 12, 2022.

Product number

The following are newly announced features on the specific models of the IBM Power 9105 machine type:

Description	Machine type	Model number	Feature number
IBM Power S1022	9105	22A	
EMEA Bulk MES Indicator	9105	22A	0004
One CSC Billing Unit	9105	22A	0010
Ten CSC Billing Units	9105	22A	0011
AIX Partition Specify	9105	22A	0265
Linux Partition Specify	9105	22A	0266
IBM i Operating System Partition Specify	9105	22A	0267
CBU Specify	9105	22A	0444
Customer Specified Placement	9105	22A	0456
Load Source Not in CEC	9105	22A	0719
Fiber Channel SAN Load Source Specify	9105	22A	0837
USB 500 GB Removable Disk Drive	9105	22A	1107
Custom Service Specify, Rochester Minn, USA	9105	22A	1140
300GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)	9105	22A	1953
600GB 10k RPM SAS SFF-2 HDD for AIX/Linux	9105	22A	1964
Primary OS - AIX	9105	22A	2146
Primary OS - Linux	9105	22A	2147
IBM i with VIOS Only System Indicator	9105	22A	2148
Factory Deconfiguration of 1-core	9105	22A	2319
1.8 M (6-ft) Extender Cable for Displays (15-pin D-shell to 15-pin D-shell)	9105	22A	4242
Rack Integration Services	9105	22A	4649
Rack Indicator- Not Factory Integrated	9105	22A	4650
Rack Indicator, Rack #1	9105	22A	4651
Rack Indicator, Rack #2	9105	22A	4652
Rack Indicator, Rack #3	9105	22A	4653
Rack Indicator, Rack #4	9105	22A	4654
Rack Indicator, Rack #5	9105	22A	4655
Rack Indicator, Rack #6	9105	22A	4656
Rack Indicator, Rack #7	9105	22A	4657
Rack Indicator, Rack #8	9105	22A	4658
Rack Indicator, Rack #9	9105	22A	4659
Rack Indicator, Rack #10	9105	22A	4660
Rack Indicator, Rack #11	9105	22A	4661
Rack Indicator, Rack #12	9105	22A	4662
Rack Indicator, Rack #13	9105	22A	4663
Rack Indicator, Rack #14	9105	22A	4664
Rack Indicator, Rack #15	9105	22A	4665
Rack Indicator, Rack #16	9105	22A	4666
Software Preload Required	9105	22A	5000

PowerVM Enterprise Edition	9105	22A	5228
PCIe2 LP 4-port 1GbE Adapter	9105	22A	5260
PCIe2 4-port 1GbE Adapter	9105	22A	5899
Power Cord 4.3m (14-ft), Drawer to IBM PDU (250V/10A)	9105	22A	6458
Power Cord 4.3m (14-ft), Drawer To OEM PDU (125V, 15A)	9105	22A	6460
Power Cord 4.3m (14-ft), Drawer to wall/OEM PDU (250V/15A) U. S.	9105	22A	6469
Power Cord 1.8m (6-ft), Drawer to wall (125V/15A)	9105	22A	6470
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/10A)	9105	22A	6471
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/16A)	9105	22A	6472
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU (250V/10A)	9105	22A	6473
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/13A)	9105	22A	6474
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/16A)	9105	22A	6475
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9105	22A	6476
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/16A)	9105	22A	6477
Power Cord 2.7 M(9-foot), To wall/OEM PDU, (250V, 16A)	9105	22A	6478
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (125V/15A or 250V/10A)	9105	22A	6488
4.3m (14-Ft) 3PH/32A 380-415V Power Cord	9105	22A	6489
4.3m (14-Ft) 1PH/63A 200-240V Power Cord	9105	22A	6491
4.3m (14-Ft) 1PH/60A (48A derated) 200-240V Power Cord	9105	22A	6492
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9105	22A	6493
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9105	22A	6494
Power Cord 2.7M (9-foot), To wall/OEM PDU, (250V, 10A)	9105	22A	6496
Power Cable - Drawer to IBM PDU, 200-240v/10A	9105	22A	6577
Power Cord 2.7M (9-foot), To wall/OEM PDU, (125V, 15A)	9105	22A	6651
4.3m (14-Ft) 3PH/16A 380-415V Power Cord	9105	22A	6653
4.3m (14-Ft) 1PH/30A (24A derated) Power Cord	9105	22A	6654
4.3m (14-Ft) 1PH/30A (24A derated) WR Power Cord	9105	22A	6655
4.3m (14-Ft) 1PH/32A Power Cord	9105	22A	6656
4.3m (14-Ft) 1PH/32A Power Cord-Australia	9105	22A	6657
4.3m (14-Ft) 1PH/30A (24A derated) Power Cord-Korea	9105	22A	6658
Power Cord 2.7M (9-foot), To wall/OEM PDU, (250V, 15A)	9105	22A	6659
Power Cord 4.3m (14-ft), Drawer to wall/OEM PDU (125V/15A)	9105	22A	6660
Power Cord 2.8m (9.2-ft), Drawer to IBM PDU, (250V/10A)	9105	22A	6665
4.3m (14-Ft) 3PH/32A 380-415V Power Cord-Australia	9105	22A	6667
Power Cord 4.3M (14-foot), Drawer to OEM PDU, (250V, 15A)	9105	22A	6669
Power Cord 2.7M (9-foot), Drawer to IBM PDU, 250V/10A	9105	22A	6671
Power Cord 2M (6.5-foot), Drawer to IBM PDU, 250V/10A	9105	22A	6672
Power Cord 2.7m (9-ft), Drawer to wall/OEM PDU, (250V/10A)	9105	22A	6680
Intelligent PDU+, 1 EIA Unit, Universal UTG0247 Connector	9105	22A	7109
Power Distribution Unit	9105	22A	7188
Power Distribution Unit (US) - 1 EIA Unit, Universal, Fixed Power Cord	9105	22A	7196
Order Routing Indicator- System Plant	9105	22A	9169
Language Group Specify - US English	9105	22A	9300
New AIX License Core Counter	9105	22A	9440
New IBM i License Core Counter	9105	22A	9441
New Red Hat License Core Counter	9105	22A	9442

New SUSE License Core Counter	9105	22A	9443
Other AIX License Core Counter	9105	22A	9444
Other Linux License Core Counter	9105	22A	9445
3rd Party Linux License Core Counter	9105	22A	9446
VIOS Core Counter	9105	22A	9447
Other License Core Counter	9105	22A	9449
Month Indicator	9105	22A	9461
Day Indicator	9105	22A	9462
Hour Indicator	9105	22A	9463
Minute Indicator	9105	22A	9464
Qty Indicator	9105	22A	9465
Countable Member Indicator	9105	22A	9466
Language Group Specify - Dutch	9105	22A	9700
Language Group Specify - French	9105	22A	9703
Language Group Specify - German	9105	22A	9704
Language Group Specify - Polish	9105	22A	9705
Language Group Specify - Norwegian	9105	22A	9706
Language Group Specify - Portuguese	9105	22A	9707
Language Group Specify - Spanish	9105	22A	9708
Language Group Specify - Italian	9105	22A	9711
Language Group Specify - Canadian French	9105	22A	9712
Language Group Specify - Japanese	9105	22A	9714
Language Group Specify - Traditional Chinese (Taiwan)	9105	22A	9715
Language Group Specify - Korean	9105	22A	9716
Language Group Specify - Turkish	9105	22A	9718
Language Group Specify - Hungarian	9105	22A	9719
Language Group Specify - Slovakian	9105	22A	9720
Language Group Specify - Russian	9105	22A	9721
Language Group Specify - Simplified Chinese (PRC)	9105	22A	9722
Language Group Specify - Czech	9105	22A	9724
Language Group Specify - Romanian	9105	22A	9725
Language Group Specify - Croatian	9105	22A	9726
Language Group Specify - Slovenian	9105	22A	9727
Language Group Specify - Brazilian Portuguese	9105	22A	9728
Language Group Specify - Thai	9105	22A	9729
Product Renovated by IBM Indicator	9105	22A	9993
10m (30.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers	9105	22A	EB2J
30m (90.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers	9105	22A	EB2K
AC Titanium Power Supply - 2000W for Server (200-240 VAC)	9105	22A	EB3N
Lift tool based on GenieLift GL-8 (standard)	9105	22A	EB3Z
10GbE Optical Transceiver SFP+ SR	9105	22A	EB46
25GbE Optical Transceiver SFP28	9105	22A	EB47
1GbE Base-T Transceiver RJ45	9105	22A	EB48
QSFP28 to SFP28 Connector	9105	22A	EB49
0.5m SFP28/25GbE copper Cable	9105	22A	EB4J
1.0m SFP28/25GbE copper Cable	9105	22A	EB4K
2.0m SFP28/25GbE copper Cable	9105	22A	EB4M
2.0m QSFP28/100GbE copper split Cable to SFP28 4x25GbE	9105	22A	EB4P
Service wedge shelf tool kit for EB3Z	9105	22A	EB4Z
QSFP+ 40GbE Base-SR4 Transceiver	9105	22A	EB57
100GbE Optical Transceiver QSFP28	9105	22A	EB59
1.0M 100GbE Copper Cable QSFP28	9105	22A	EB5K
1.5M 100GbE Copper Cable QSFP28	9105	22A	EB5L
2.0M 100GbE Copper Cable QSFP28	9105	22A	EB5M
3M 100GbE Optical Cable QSFP28 (AOC)	9105	22A	EB5R
5M 100GbE Optical Cable QSFP28 (AOC)	9105	22A	EB5S
10M 100GbE Optical Cable QSFP28 (AOC)	9105	22A	EB5T
15M 100GbE Optical Cable QSFP28 (AOC)	9105	22A	EB5U
20M 100GbE Optical Cable QSFP28 (AOC)	9105	22A	EB5V
30M 100GbE Optical Cable QSFP28 (AOC)	9105	22A	EB5W
50M 100GbE Optical Cable QSFP28 (AOC)	9105	22A	EB5X
IBM i 7.3 Indicator	9105	22A	EB73
IBM i 7.4 Indicator	9105	22A	EB74
IBM i 7.5 Indicator	9105	22A	EB75
PCIe3 LP 2-Port 10Gb NIC&ROCE SR/Cu Adapter	9105	22A	EC2R
PCIe3 2-Port 10Gb NIC&ROCE SR/Cu Adapter	9105	22A	EC2S
PCIe3 LP 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter	9105	22A	EC2T
PCIe3 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter	9105	22A	EC2U
PCIe3 x8 LP 3.2 TB NVME Flash adapter for AIX/			

Linux	9105	22A	EC5C
PCIe3 x8 LP 6.4 TB NVMe Flash adapter for AIX/Linux	9105	22A	EC5E
PCIe3 x8 LP 1.6 TB NVMe Flash Adapter for AIX/Linux	9105	22A	EC5G
Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for AIX/Linux	9105	22A	EC5V
Mainstream 800 GB SSD PCIe3 NVMe U.2 module for AIX/Linux	9105	22A	EC5X
PCIe4 LP 2-port 100Gb ROCE EN LP adapter	9105	22A	EC67
PCIe2 LP 2-Port USB 3.0 Adapter	9105	22A	EC6J
PCIe2 2-Port USB 3.0 Adapter	9105	22A	EC6K
PCIe4 LP 1.6TB NVMe Flash Adapter x8 for AIX/Linux	9105	22A	EC7A
PCIe4 LP 3.2TB NVMe Flash Adapter x8 for AIX/Linux	9105	22A	EC7C
PCIe4 LP 6.4TB NVMe Flash Adapter x8 for AIX/Linux	9105	22A	EC7E
800GB Mainstream NVMe U.2 SSD 4k for AIX/Linux	9105	22A	EC7T
SAS X Cable 3m - HD Narrow 6Gb 2-Adapters to Enclosure	9105	22A	ECBJ
SAS X Cable 6m - HD Narrow 6Gb 2-Adapters to Enclosure	9105	22A	ECBK
SAS Y0 Cable 1.5m - HD Narrow 6Gb Adapter to Enclosure	9105	22A	ECBT
SAS Y0 Cable 3m - HD Narrow 6Gb Adapter to Enclosure	9105	22A	ECBU
SAS Y0 Cable 6m - HD Narrow 6Gb Adapter to Enclosure	9105	22A	ECBV
SAS Y0 Cable 10m - HD Narrow 6Gb Adapter to Enclosure	9105	22A	ECBW
SAS AE1 Cable 4m - HD Narrow 6Gb Adapter to Enclosure	9105	22A	ECBY
SAS YE1 Cable 3m - HD Narrow 6Gb Adapter to Enclosure	9105	22A	ECBZ
System Port Converter Cable for UPS	9105	22A	ECCF
3M Copper CXP Cable Pair for PCIe3 Expansion Drawer	9105	22A	ECCS
3.0M SAS X12 Cable (Two Adapter to Enclosure)	9105	22A	ECDJ
4.5M SAS X12 Active Optical Cable (Two Adapter to Enclosure)	9105	22A	ECDK
10M SAS X12 Active Optical Cable (Two Adapter to Enclosure)	9105	22A	ECDL
1.5M SAS Y012 Cable (Adapter to Enclosure)	9105	22A	ECDT
3.0M SAS Y012 Cable (Adapter to Enclosure)	9105	22A	ECDU
4.5M SAS Y012 Active Optical Cable (Adapter to Enclosure)	9105	22A	ECDV
10M SAS Y012 Active Optical Cable (Adapter to Enclosure)	9105	22A	ECDW
0.6M SAS AA12 Cable (Adapter to Adapter)	9105	22A	ECE0
3.0M SAS AA12 Cable	9105	22A	ECE3
4.5M SAS AA12 Active Optical Cable (Adapter to Adapter)	9105	22A	ECE4
4.3m (14-Ft) PDU to wall 3PH/24A 200-240V Delta-wired Power Cord	9105	22A	ECJ5
4.3m (14-Ft) PDU to wall 3PH/40A 200-240V Power Cord	9105	22A	ECJ6
4.3m (14-Ft) PDU to wall 3PH/48A 200-240V Delta-wired Power Cord	9105	22A	ECJ7
High Function 9xC19 Single-Phase or Three-Phase wye PDU plus	9105	22A	ECJJ
High Function 9xC19 PDU plus 3-Phase Delta	9105	22A	ECJL
High Function 12xC13 Single-Phase or Three-Phase wye PDU plus	9105	22A	ECJN
High Function 12xC13 PDU plus 3-Phase Delta	9105	22A	ECJQ
Custom Service Specify, Mexico	9105	22A	ECSM
Custom Service Specify, Poughkeepsie, USA	9105	22A	ECSP
Optical wrap Plug	9105	22A	ECW0
SAP HANA TRACKING FEATURE	9105	22A	EHKV
Power10 Healthcare Solution Edition indicator	9105	22A	EHM1
Boot Drive / Load Source in EXP24SX Specify (in #ESLS or #ELLS)	9105	22A	EHR2
SSD Placement Indicator - #ESLS/#ELLS	9105	22A	EHS2
PCIe3 RAID SAS Adapter Quad-port 6Gb x8	9105	22A	EJ0J

PCIe3 12GB Cache RAID SAS Adapter Quad-port 6Gb x8	9105	22A	EJ0L
PCIe3 LP RAID SAS Adapter Quad-Port 6Gb x8	9105	22A	EJ0M
PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8	9105	22A	EJ10
PCIe3 LP SAS Tape/DVD Adapter Quad-port 6Gb x8	9105	22A	EJ11
PCIe3 12GB Cache RAID PLUS SAS Adapter Quad-port 6Gb x8	9105	22A	EJ14
PCIe x16 to CXP Optical or CU converter Adapter for PCIe3 Expansion Drawer	9105	22A	EJ1R
Storage Backplane with four NVMe U.2 drive slots	9105	22A	EJ1X
PCIe x16 to CXP Converter Card, Supports optical cables	9105	22A	EJ24
PCIe3 Crypto Coprocessor BSC-Gen3 4767	9105	22A	EJ33
PCIe3 Crypto Coprocessor BSC-Gen3 4769	9105	22A	EJ37
Non-paired Indicator EJ14 PCIe SAS RAID+ Adapter	9105	22A	EJRL
Non-paired Indicator EJ0L PCIe SAS RAID Adapter	9105	22A	EJRU
Front IBM Bezel for 8 NVMe-bays Backplane Rack-Mount	9105	22A	EJUS
Front OEM Bezel for 8 NVMe-bays Backplane Rack-Mount	9105	22A	EJUT
Specify Mode-1 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP24SX #ESLS/ELLS	9105	22A	EJW1
Specify Mode-1 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9105	22A	EJW2
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJW3
Specify Mode-2 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJW4
Specify Mode-4 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJW5
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9105	22A	EJW6
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)Y012 for EXP24SX #ESLS/ELLS	9105	22A	EJW7
Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)Y012 for EXP24SX #ESLS/ELLS	9105	22A	EJWA
Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJWB
Specify Mode-4 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJWC
Specify Mode-4 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJWD
Specify Mode-4 & (3)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJWE
Specify Mode-1 & (2)EJ14 & (2)Y012 for EXP24SX #ESLS/ELLS	9105	22A	EJWF
Specify Mode-2 & (2)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJWG
Specify Mode-2 & (2)EJ14 & (1)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJWH
Specify Mode-2 & (4)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS	9105	22A	EJWJ
Specify Mode-1 & CEC SAS port Controller EJ1G/EL67 & (1)Y012 for EXP24SX #ESLS/ELLS	9105	22A	EJWU
300GB 15k RPM SAS SFF-2 Disk Drive (Linux)	9105	22A	EL1P
600GB 10k RPM SAS SFF-2 Disk Drive (Linux)	9105	22A	EL1Q
PDU Access Cord 0.38m	9105	22A	ELC0
4.3m (14-Ft) PDU to wall 24A 200-240V Power Cord North America	9105	22A	ELC1
4.3m (14-Ft) PDU to wall 3PH/24A 415V Power Cord North America	9105	22A	ELC2
Power Cable - Drawer to IBM PDU (250V/10A)	9105	22A	ELC5
600GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	9105	22A	ELEV
1.2TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	9105	22A	ELF3
1.8TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096	9105	22A	ELFT
IBM Power Private Cloud Rack Solution Indicator	9105	22A	ELG2
512GB Memory Bundle for Power10 Healthcare Solution Edition	9105	22A	EM68
32GB (2x16GB) DDIMMs, 3200 MHZ, 8GBIT DDR4 Memory	9105	22A	EM6N
64GB (2x32GB) DDIMMs, 3200 MHZ, 8GBIT DDR4 Memory	9105	22A	EM6W
128GB (2x64GB) DDIMMs, 3200 MHZ, 16GBIT DDR4			

Memory	9105	22A	EM6X
256GB (2x128GB) DDIMMs, 2666 MHZ, 16GBIT DDR4			
Memory	9105	22A	EM6Y
Active Memory Mirroring (AMM)	9105	22A	EM8G
PCIe Gen3 I/O Expansion Drawer	9105	22A	EMX0
AC Power Supply Conduit for PCIe3 Expansion			
Drawer	9105	22A	EMXA
PCIe3 6-Slot Fanout Module for PCIe3 Expansion			
Drawer	9105	22A	EMXF
PCIe3 6-Slot Fanout Module for PCIe3 Expansion			
Drawer	9105	22A	EMXG
PCIe3 6-Slot Fanout Module for PCIe3 Expansion			
Drawer	9105	22A	EMXH
1m (3.3-ft), 10Gb E'Net Cable SFP+ Act Twinax			
Copper	9105	22A	EN01
3m (9.8-ft), 10Gb E'Net Cable SFP+ Act Twinax			
Copper	9105	22A	EN02
5m (16.4-ft), 10Gb E'Net Cable SFP+ Act Twinax			
Copper	9105	22A	EN03
PCIe2 4-Port (10Gb+1GbE) SR+RJ45 Adapter	9105	22A	EN0S
PCIe2 LP 4-Port (10Gb+1GbE) SR+RJ45 Adapter	9105	22A	EN0T
PCIe2 4-port (10Gb+1GbE) Copper SFP+RJ45 Adapter	9105	22A	EN0U
PCIe2 LP 4-port (10Gb+1GbE) Copper SFP+RJ45			
Adapter	9105	22A	EN0V
PCIe2 2-port 10/1GbE BaseT RJ45 Adapter	9105	22A	EN0W
PCIe2 LP 2-port 10/1GbE BaseT RJ45 Adapter	9105	22A	EN0X
PCIe3 32Gb 2-port Fibre Channel Adapter	9105	22A	EN1A
PCIe3 LP 32Gb 2-port Fibre Channel Adapter	9105	22A	EN1B
PCIe3 16Gb 4-port Fibre Channel Adapter	9105	22A	EN1C
PCIe3 LP 16Gb 4-port Fibre Channel Adapter	9105	22A	EN1D
PCIe3 16Gb 4-port Fibre Channel Adapter	9105	22A	EN1E
PCIe3 LP 16Gb 4-port Fibre Channel Adapter	9105	22A	EN1F
PCIe3 2-Port 16Gb Fibre Channel Adapter	9105	22A	EN1G
PCIe3 LP 2-Port 16Gb Fibre Channel Adapter	9105	22A	EN1H
PCIe4 32Gb 2-port Optical Fibre Channel Adapter	9105	22A	EN1J
PCIe4 LP 32Gb 2-port Optical Fibre Channel			
Adapter	9105	22A	EN1K
PCIe3 16Gb 2-port Fibre Channel Adapter	9105	22A	EN2A
PCIe3 LP 16Gb 2-port Fibre Channel Adapter	9105	22A	EN2B
Power Enterprise Pools 2.0 Enablement	9105	22A	EP20
Deactivation of LPM (Live Partition Mobility)	9105	22A	EPA0
One CUoD Static Processor Core Activation for			
EPG8	9105	22A	EPF8
One CUoD Static Processor Core Activation for			
EPG9	9105	22A	EPF9
One CUoD Static Processor Core Activation for			
EPGA	9105	22A	EPFA
16-core Typical 2.75 to 4.0 Ghz (max) Power10			
Processor	9105	22A	EPG8
12-core Typical 2.90 to 4.0 Ghz (max) Power10			
Processor	9105	22A	EPG9
20-core Typical 2.45 to 3.90 Ghz (max) Power10			
Processor	9105	22A	EPGA
Horizontal PDU Mounting Hardware	9105	22A	EPTH
High Function 9xC19 PDU: Switched, Monitoring	9105	22A	EPTJ
High Function 9xC19 PDU 3-Phase: Switched,			
Monitoring	9105	22A	EPTL
High Function 12xC13 PDU: Switched, Monitoring	9105	22A	EPTN
High Function 12xC13 PDU 3-Phase: Switched,			
Monitoring	9105	22A	EPTQ
Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for			
AIX/Linux	9105	22A	ES1E
Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for			
AIX/Linux	9105	22A	ES1G
Enterprise 1.6 TB SSD PCIe4 NVMe U.2 module for			
AIX/Linux	9105	22A	ES3B
Enterprise 3.2 TB SSD PCIe4 NVMe U.2 module for			
AIX/Linux	9105	22A	ES3D
Enterprise 6.4 TB SSD PCIe4 NVMe U.2 module for			
AIX/Linux	9105	22A	ES3F

387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ES94
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ESB2
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ESB6
387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESBA
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESBG
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESBL
S&H - No Charge	9105	22A	ESC0
S&H-a	9105	22A	ESC5
Virtual Capacity Expedited Shipment	9105	22A	ESCT
iSCSI SAN Load Source Specify for AIX	9105	22A	ESCZ
600GB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux	9105	22A	ESEV
1.2TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux	9105	22A	ESF3
1.8TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux	9105	22A	ESFT
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ESGV
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ESGZ
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJ0
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJ2
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJ4
7.45TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJ6
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJJ
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJL
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJN
7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESJQ
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ESK1
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ESK3
387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESK8
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESKC
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESKG
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESKK
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESKP
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESKT
7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESKX
Specify AC Power Supply for EXP12SX/EXP24SX			
Storage Enclosure	9105	22A	ESLA
EXP24SX SAS Storage Enclosure	9105	22A	ESLS
931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESMB
1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESMF
3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESMK
7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESMV
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESNA
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ESNE
300GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	9105	22A	ESNM
600GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)	9105	22A	ESNR
300GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (Linux)	9105	22A	ESRM
600GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (Linux)	9105	22A	ESRR
AIX Update Access Key (UAK)	9105	22A	ESWK
387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ETK1
775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux	9105	22A	ETK3
387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ETK8
775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ETKC
1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux	9105	22A	ETKG
1TB Removable Disk Drive Cartridge	9105	22A	EU01
RDX 320 GB Removable Disk Drive	9105	22A	EU08
Operator Panel LCD Display	9105	22A	EU0K
1.5TB Removable Disk Drive Cartridge	9105	22A	EU15
Cable Ties & Labels	9105	22A	EU19
Order Placed Indicator	9105	22A	EU29
2TB Removable Disk Drive Cartridge (RDX)	9105	22A	EU2T
RDX USB External Docking Station	9105	22A	EUA4

Note: Feature EUA4 is not supported in Armenia, Azerbaijan, China, India, Japan, Kazakhstan, Kyrgyzstan, Mexico, Saudi Arabia, Taiwan, Turkmenistan, and Uzbekistan.

Standalone USB DVD drive w/cable	9105	22A	EUA5
1 core Base Processor Activation (Pools 2.0) for EPG8 - Any OS	9105	22A	EUCA

1 core Base Processor Activation (Pools 2.0) for EPG9 - Any OS	9105	22A	EUCB
1 core Base Processor Activation (Pools 2.0) for EPGA - Any OS	9105	22A	EUCC
1 core Base Processor Activation (Pools 2.0) for EPG9 - Any O/S (Conv from EPF9)	9105	22A	EUCH
1 core Base Processor Activation (Pools 2.0) for EPG8 - Any O/S (Conv from EPF8)	9105	22A	EUCG
1 core Base Processor Activation (Pools 2.0) for EPGA - Any O/S (Conv from EPFA)	9105	22A	EUCJ
Enable Virtual Serial Number	9105	22A	EVSN
BP Post-Sale Services: 1 Day	9105	22A	SVBP
IBM Systems Lab Services Post-Sale Services: 1 Day			
Other IBM Post-Sale Services: 1 Day	9105	22A	SVCS
			SVNN

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

Planned Availability Date July 22, 2022

New Feature

Description	Machine type	Model number	Feature number
Rack Content Specify	9105-22A, 9786-22H	2E1A unit 7965	S42 ER39

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 clients:

Feature conversions for 9105-22A adapter features:

From FC:	To FC:	Return parts
EJ1R - 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 9105-22A processor features:

From FC:	To FC:	Return parts
EPF8 - One CUoD Static Processor Core Activation for EPG8	EUCG - 1 core Base Processor Activation (Pools 2.0) for EPG8 - Any OS (Conv from EPF8)	No
EPF9 - One CUoD Static Processor Core Activation for EPG9	EUCH - 1 core Base Processor Activation (Pools 2.0) for EPG9 - Any OS (Conv from EPF9)	No
EPFA - One CUoD Static Processor Core Activation for EPGA	EUCJ - 1 core Base Processor Activation (Pools 2.0) for EPGA - Any OS	No

Feature conversions for 9105-22A 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

No publications are shipped with the announced product.

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

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

National language support

Not applicable

Services

IBM Systems Lab Services

Systems Lab Services offers infrastructure services to help build hybrid cloud and enterprise IT solutions. From servers to storage systems and software, Systems Lab Services can help deploy the building blocks of a next-generation IT infrastructure to empower a client's business. Systems Lab Services consultants can perform infrastructure services for clients online or onsite, offering deep technical expertise, valuable tools, and successful methodologies. Systems Lab Services is designed to help clients solve business challenges, gain new skills, and apply best practices.

Systems Lab Services offers a wide range of infrastructure services for IBM Power servers, IBM Storage systems, IBM Z^(R), and IBM LinuxONE. Systems Lab Services has a global presence and can deploy experienced consultants online or onsite around the world.

For assistance, contact Systems Lab Services at ibmsls@us.ibm.com.

To learn more, see the [IBM Systems Lab Services](#) website.

IBM ConsultingTM

As transformation continues across every industry, businesses need a single partner to map their enterprise-wide business strategy and technology infrastructure. IBM Consulting is the business partner to help accelerate change across an organization. IBM specialists can help businesses succeed through finding collaborative ways of working that forge connections across people, technologies, and partner ecosystems. IBM Consulting brings together the business expertise and an ecosystem of technologies that help solve some of the biggest problems faced by organizations. With methods that get results faster, an integrated approach that is grounded in an open and flexible hybrid cloud architecture, and incorporating technology from IBM Research^(R) and IBM Watson^(R) AI, IBM Consulting enables

businesses to lead change with confidence and deliver continuous improvement across a business and its bottom line.

For additional information, see the [IBM Consulting](#) website.

IBM Technology Support Services (TSS)

Get preventive maintenance, onsite and remote support, and gain actionable insights into critical business applications and IT systems. Speed developer innovation with support for over 240 open-source packages. Leverage powerful IBM analytics and AI-enabled tools to enable client teams to manage IT problems before they become emergencies.

TSS offers extensive IT maintenance and support services that cover more than one niche of a client's environment. TSS covers products from IBM and OEMs, including servers, storage, network, appliances, and software, to help clients ensure high availability across their data center and hybrid cloud environment.

For details on available services, see the [Technology support for hybrid cloud environments](#) website.

IBM Expert Labs

Expert Labs can help clients accelerate their projects and optimize value by leveraging their deep technical skills and knowledge. With more than 20 years of industry experience, these specialists know how to overcome the biggest challenges to deliver business results that can have an immediate impact.

Expert Labs' deep alignment with IBM product development allows for a strategic advantage as they are often the first in line to get access to new products, features, and early visibility into roadmaps. This connection with the development enables them to deliver First of a Kind implementations to address unique needs or expand a client's business with a flexible approach that works best for their organization.

For additional information, see the [IBM Expert Labs](#) website.

IBM Security^(R) Expert Labs

With extensive consultative expertise on IBM Security software solutions, Security Expert Labs helps clients and partners modernize the security of their applications, data, and workforce. With an extensive portfolio of consulting and learning services, Expert Labs provides project-based and premier support service subscriptions.

These services can help clients deploy and integrate IBM Security software, extend their team resources, and help guide and accelerate successful hybrid cloud solutions, including critical strategies such as zero trust. Remote and on-premises software deployment assistance is available for IBM Cloud Pak^(R) for Security, IBM Security QRadar^(R)/QRoC, IBM Security SOAR/Resilient^(R), IBM i2^(R), IBM Security Verify, IBM Security Guardium^(R), and IBM Security MaaS360^(R).

For more information, contact Security Expert Labs at sel@us.ibm.com.

For additional information, see the [IBM Security Expert Labs](#) website.

IBM support

For installation and technical support information, see the [IBM Support Portal](#).

Additional support

IBM Client Engineering for Systems

Client Engineering for Systems is a framework for accelerating digital transformation. It helps you generate innovative ideas and equips you with the practices, technologies, and expertise to turn those ideas into business value in weeks. When you work with Client Engineering for Systems, you bring pain points into focus. You empower your team to take manageable risks, adopt leading technologies, speed up solution development, and measure the value of everything you do. Client Engineering for Systems has experts and services to address a broad array of use cases, including capabilities for business transformation, hybrid cloud, analytics and AI, infrastructure systems, security, and more. Contact Client Engineering at sysgarage@ibm.com.

Technical information

Specified operating environment

Physical specifications

- Width¹: 482 mm (18.97 in.)
- Depth²: 813 mm (32 in.)
- Height: 86.5 mm (3.4 in.)
- Weight: 32.20 kg (71 lb)

1. The width is measured to the outside edges of the rack-mount bezels. The width is 446 mm (17.6 in.) for the main chassis, which fits in between a 482.6 mm (19 in.) rack mounting flanges.
2. The cable management arm with the maximum cable bundle adds 241 mm (9.5 in.) to the depth. Feature ECRK is required for the 7965-S42 rack.

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

Operating environment

Electrical characteristics

- AC rated voltage and frequency²: 200--240 V AC at 50 or 60 Hz plus or minus 3 Hz
- Thermal output (maximum)³: 7643 BTU/hr
- Maximum power consumption³: 2240 W
- Maximum kVA⁴: 2.31 kVA
- Phase: Single

1. Redundancy is supported. The Power S1022 has a maximum of two power supplies. There are no specific plugging rules or plugging sequence when you connect the power supplies to the rack PDUs. All the power supplies feed a common DC bus.

2. The power supplies automatically accept any voltage with the published, rated-voltage range. If multiple power supplies are installed and operating, the power supplies draw approximately equal current from the utility (electrical supply) and provide approximately equal current to the load.

3. Power draw and heat load vary greatly by configuration. When you plan for an electrical system, it is important to use the maximum values. However, when you plan for heat load, you can use the IBM Systems Energy Estimator to obtain a heat output estimate based on a specific configuration. For more information, see [The IBM Systems Energy Estimator website](#).

4. To calculate the amperage, multiply the kVA by 1,000 and divide that number by the operating voltage.

Environment (operating)¹

- ASHRAE class; allowable A3 (fourth edition)
- Airflow direction; recommended Front-to-back
- Temperature: Recommended 18.0°C--27.0°C (64.4°F--80.6°F); allowable 5.0°C--40.0°C (41.0°F--104.0°F)
- Low end moisture: Recommended 9.0°C (15.8°F) dew point; allowable -12.0°C (10.4°F) dew point and 8% relative humidity
- High-end moisture: Recommended 60% relative humidity and 15°C (59°F) dew point; allowable 85% relative humidity and 24.0°C (75.2°F) dew point
- Maximum altitude: 3,050 m (10,000 ft)

Allowable environment (nonoperating)⁵

- Temperature: Recommended 5°C--45°C (41°F--113°F)
- Relative humidity: Recommended 8% to 85%
- Maximum dew point: Recommended 27.0°C (80.6°F)

1. IBM provides the recommended operating environment as the long-term operating environment that can result in the greatest reliability, energy efficiency, and reliability. The allowable operating environment represents where the equipment is tested to verify functionality. Due to the stresses that operating in the allowable envelope can place on the equipment, these envelopes must be used for short-term operation, not continuous operation. There are a very limited number of configurations that must not operate at the upper bound of the A3 allowable range. For more information, consult your IBM technical specialist.

2. Must derate the maximum allowable temperature 1°C (1.8°F) per 175 m (574 ft) above 900 m (2,953 ft) up to a maximum allowable elevation of 3,050 m (10,000 ft).

3. The minimum humidity level is the larger absolute humidity of the -12°C (10.4°F) dew point and the 8% relative humidity. These levels intersect at approximately 25°C (77°F). Below this intersection, the dew point (-12°C) represents the minimum moisture level, while above it, the relative humidity (8%) is the minimum. For the upper moisture limit, the limit is the minimum absolute humidity of the dew point and relative humidity that is stated.

4. The following minimum requirements apply to data centers that are operated at low relative humidity:

- Data centers that do not have ESD floors and where people are allowed to wear non-ESD shoes might want to consider increasing humidity given that the risk of generating 8 kV increases slightly at 8% relative humidity, when compared to 25% relative humidity.
- All mobile furnishings and equipment must be made of conductive or static dissipative materials and be bonded to ground.
- During maintenance on any hardware, a properly functioning and grounded wrist strap must be used by any personnel who comes in contact with information technology (IT) equipment.

5. Equipment that is removed from the original shipping container and is installed, but is powered down. The allowable nonoperating environment is provided to define the environmental range that an unpowered system can experience short term without being damaged.

Electromagnetic compatibility compliance: CISPR 22; CISPR 32; CISPR 24; CISPR 35; FCC, CFR 47, Part 15 (US); VCCI (Japan); EMC Directive (EEA); ICES-003 (Canada); ACMA (Australia, New Zealand); CNS 13438 (Taiwan); Radio Waves Act

(Korea); Commodity Inspection Law (China); QCVN 118 (Vietnam); MoCI (Saudi Arabia); SI 961 (Israel); EAC (EAEU).

Safety compliance: This product was designed, tested, manufactured, and certified for safe operation. It complies with IEC 60950-1 and/or IEC 62368-1 and where required, to relevant national differences/deviations (ND) to these IEC base standards. This includes, but is not limited to: EN (European Norms including all Amendments under the Low Voltage Directive), UL/CSA (North America bi-national harmonized and marked per accredited NRTL agency listings), and other such derivative certifications according to corporate determinations and latest regional publication compliance standardized requirements.

See the *Installation Planning Guide* in [IBM Documentation](#) for additional detail.

Homologation

This product is not certified for direct connection by any means whatsoever to interfaces of public telecommunications networks. Certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Hardware requirements

Power S1022 system configuration

The minimum Power S1022 initial order must include a processor module, two 16 GB DIMMs (one feature EM6N 32 GB (2 x 16 GB) DDIMM), two power supplies and line cords, an operating system indicator, a cover set indicator, and a Language Group Specify. Also, it must include one of these storage options and one of these network options:

Storage options:

- For boot from NVMe: One NVMe drive slot and one NVMe drive or one PCIe NVMe add in adapter.
- For boot from SAN: Internal HDD or SSD and RAID card are **not** required if feature 0837 (boot from SAN) is selected. A Fibre Channel adapter must be ordered if feature 0837 is selected.

Network options:

- One PCIe2 4-port 1 Gb Ethernet adapter
- One of the supported 10 Gb Ethernet adapters

AIX or Linux is the primary operating system. The minimum defined initial order configuration is as follows:

System Feature Codes	Feature Code	Description	Default	Minimum Quantity	Notes
Op-Panel	EU0K	Operator Panel LCD Display		1	Optional with AIX/Linux. Always default Qty. 1, but can be deselected for AIX/Linux.
Virtualization Engine	5228	PowerVM Enterprise Edition	1	1	Must select one option.
or					

System Feature Codes	Feature Code	Description	Default	Minimum Quantity	Notes
	EPA0	Deactivation of LPM (Live Partition Mobility)		1	
Processor Modules	EPG9	12-core Typical 2.90 to 4.0 Ghz (max) Power10 Processor		1	Must select Processor Module option.
	or				
	EPG8	16-core Typical 2.75 to 4.0 Ghz (max) Power10 Processor		2	
	or				
	EPGA	20-core Typical 2.45 to 3.90 Ghz (max) Power10 Processor		2	
Processor Module Activations	EPF9	One CUoD Static Processor Core Activation for EPG9		6	Minimum of 50% of CUoD Static processor core activations need to be ordered.
	or				
	EPF8	One CUoD Static Processor Core Activation for EPG8		16	
	or				
	EPFA	One CUoD Static Processor Core Activation for EPGA		20	
	or				
	EUCB	1 core Base Processor Activation (Pools 2.0) for EPG9 - Any OS		from 1 to 24	Requires Pools 2.0 feature EP20 Power Enterprise Pools 2.0 Enablement.
	or				
	EUCA	1 core Base Processor Activation (Pools		from 1 to 32	

System Feature Codes	Feature Code	Description	Default	Minimum Quantity	Notes
		2.0) for EPG8 - Any OS			
	or				
	EUCC	1 core Base Processor Activation (Pools 2.0) for EPGA - Any OS		from 1 to 40	
Memory	EM6N	32GB (2x16GB) DDIMMs, 3200 MHz, 8GBIT DDR4 Memory		1	Minimum 2 DIMMs = 1 DIMM feature.
	or				
	EM6W	64GB (2x32GB) DDIMMs, 3200 MHz, 8GBIT DDR4 Memory		1	
	or				
	EM6X	128GB (2x64GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory		1	
	or				
	EM6Y	256GB (2x128GB) DDIMMs, 2666 MHz, 16GBIT DDR4 Memory		1	
Active Memory Mirroring	EM8G	Active Memory Mirroring (AMM)	0	0	Optional feature. Max. Qty. 1 per system. Memory Mirroring requires a minimum of 8 DIMMs (4 features DIMM).
Storage Backplane	EJ1X	Storage Backplane with four NVMe U.2 drive slots		1	Must order Qty. 1 NVMe backplane feature except when

System Feature Codes	Feature Code	Description	Default	Minimum Quantity	Notes
					#0837 or #ESCZ (iSCSI boot) is on the order or when NVMe PCIe add-in adapter card is used as the Load Source. Mixing NVMe devices is allowed on each backplane.
Bezels	EJUS	Front IBM Bezel for 8 NVMe-bays Backplane Rack-Mount		1	When no NVMe backplane is ordered, default #EJUS.
	or EJUT	Front OEM Bezel for 8 NVMe-bays Backplane Rack-Mount		1	
NVMe Devices	EC7T	800GB Mainstream NVMe U.2 SSD 4k for AIX/Linux	2	0	For AIX/Linux, default is Qty. 2. It is allowed to be changed to any other quantity from 0 to 8.
Required LAN adapters	EC2T	PCIe3 LP 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter	1	1	Qty. 1 of these LAN features required on all Initial orders. Default Adapter: feature EC2T.
	or EN0X	PCIe2 LP 2-port 10/1GbE BaseT RJ45 Adapter		1	
Power Supply	EB3N	AC Power Supply - 2000W for Server	2	2	Each initial order must have all power

System Feature Codes	Feature Code	Description	Default	Minimum Quantity	Notes
		(200-240 VAC)			supplies present, power supplies cannot be added later on. Only 200--240V power cords can be used.
Power Cables	6458	Power Cord 4.3m (14-ft), Drawer to IBM PDU (250V/10A)	2	2	Qty. 2 required.
Language Group	9300	Language Group Specify - US English	1	1	Language Specify code is required.
Primary Operating	2146	Primary OS -AIX		1	Must select one option.
	or 2147	Primary OS -Linux		1	

1. The racking approach for the initial order can be a MTM 7965-S42.

HMC machine code

If the system is ordered with 1020 firmware level, or higher, and is capable to be HMC managed, then the managing HMC must be installed with HMC 10.1.1020.0, or higher.

This level only supports hardware appliance types 7063, virtual appliances (vHMC) on x86, or PowerVM. The 7042 hardware appliance is not supported.

An HMC is required to manage the Power S1022 server implementing partitioning. Multiple IBM Power8, Power9 and Power10 processor-based servers can be supported by a single HMC with version 10.

Planned HMC hardware and software support:

- Hardware Appliance: 7063-CR1, 7063-CR2
- vHMC on x86
- vHMC PowerVM-based LPAR

If you are attaching an HMC to a new server or adding function to an existing server that requires a firmware update, the HMC machine code may need to be updated because HMC code must always be equal to or higher than the managed server's firmware. Access to firmware and 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 at its discretion.

To determine the HMC machine code level required for the firmware level on any server, go to the following web page to access the Fix Level Recommendation Tool (FLRT) on or after the planned availability date for this product. FLRT will identify the

correct HMC machine code for the selected system firmware level; see the website [Fix Level Recommendation Tool](#).

If a single HMC is attached to multiple servers, the HMC machine code level must be updated to be at or higher than the server with the most recent firmware level. All prior levels of server firmware are supported with the latest HMC machine code level.

For clients installing systems higher than the EIA 29 position (location of the rail that supports the rack-mounted server) in any IBM or non-IBM rack, acquire approved tools outlined in the server specifications section at [IBM Documentation](#).

In situations where IBM service is required and the recommended tools are not available, there could be delays in repair actions.

Software requirements

- Red Hat Enterprise Linux 9.0, for Power LE, or later
- Red Hat Enterprise Linux 8.4, for Power LE, or later
- SUSE Linux Enterprise Server 15 Service Pack 3, or later
- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 15 Service Pack 3, or later
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8.4 for Power LE, or later
- Red Hat OpenShift Container Platform 4.9, or later

Please review the Linux alert page for any known Linux issues or limitations [Linux on IBM- Readme first issues](#) website.

If installing IBM i:

- VIOS is required, the IBM i partitions must be set to "restricted I/O" mode. There are limitations to the maximum size of the partition. Up to four cores (real or virtual) per IBM i partition are supported. Multiple IBM i partitions can be created and run concurrently, and each individual partition can have up to four cores.
- The IBM i operating system levels supported are:
 - IBM i 7.5, or later
 - IBM i 7.4 TR6, or later
 - IBM i 7.3 TR12, or later

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

- AIX Version 7.3 with the 7300-00 Technology Level and Service Pack 7300-00-02-2220, or later
- AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-04-2220, or later
- AIX Version 7.2 with the 7200-04 Technology Level and Service Pack 7200-04-06-2220, or later (planned availability September 16, 2022)

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

- AIX Version 7.3 with the 7300-00 Technology Level and service pack 7300-00-01-2148, or later
- 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 7200-04-02-2016, or later
- AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-06-2016, or later

If installing VIOS:

- VIOS 3.1.3.21

Limitations

There is no physical system port on the scale-out Power10 servers.

Planning information

Cable orders

No cables required.

Security, auditability, and control

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

Terms and conditions

Volume orders

Contact your IBM representative.

Products - terms and conditions

Warranty period

Warranty and additional coverage options:	Coverage summary⁽¹⁾:
Warranty Period:	3 years
Service Level:	IBM CRU & On-Site, 9x5 Next Business Day
Service Upgrade Options :	
Warranty Service Upgrade	IBM On-Site Repair, 9x5 Same Day ⁽²⁾ and 24x7 Same Day options
Maintenance Services (Post-Warranty):	IBM On-Site Repair, Next Business Day and Same Day options
IBM Hardware Maintenance Services - committed maintenance ⁽³⁾ :	Y

⁽¹⁾ See complete coverage details below.

⁽²⁾ Offered in US and EMEA only.

⁽³⁾ Not offered in the US.

To obtain copies of the IBM Statement of Limited Warranty, contact your reseller or IBM.

An IBM part or feature installed during the initial installation of an IBM machine is subject to the full warranty period specified by IBM. An IBM part or feature that replaces a previously installed part or feature assumes the remainder of the warranty period for the replaced part or feature. An IBM part or feature added to a machine without replacing a previously installed part or feature is subject to a full warranty. Unless specified otherwise, the warranty period, type of warranty service,

and service level of a part or feature are the same as those for the machine in which it is installed.

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 SDD 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.

The IBM warranty covers feature number EB4Z. For warranty terms associated with feature number EB3Z and the Lift tool based on GenieLift GL-8, see the separate warranty terms provided by Genie found in the Genie Operator's Manual at the [Genie](#) website.

For clients installing systems higher than the EIA 29 position (location of the rail that supports the rack-mounted server) in any IBM or non-IBM rack, acquire approved tools outlined in the server specifications section at [IBM Documentation](#). In situations where IBM service is required and the recommended tools are not available, there could be delays in repair actions.

Extended Warranty Service

Extended Warranty Service is not applicable.

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 (mandatory) CRU

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

The following parts have been designated as Tier 1 CRUs:

- Bezel
- Service Cover
- Op Panel
- Op Panel -- LCD
- Op Panel -- LCD Cable
- Fan

- Fan Card Signal Cable
- Front USB Cable
- NVMe drive
- NVMe Filler
- DDIMM Cover for Retention
- DDIMM Filler
- Time of Day Battery
- TPM Card
- Processor VRM
- Processor Heatsink
- PCIe Adapter
- Power Supply
- Power Distribution Signal Cable

Tier 2 (optional) CRU

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

Based upon availability, CRUs will be shipped for next-business-day (NBD) delivery. IBM specifies, in the materials shipped with a replacement CRU, whether a defective CRU must be returned to IBM. When return is required, return instructions and a container are shipped with the replacement CRU. You may be charged for the replacement CRU if IBM does not receive the defective CRU within 15 days of your receipt of the replacement.

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:

- 9 hours per day, Monday through Friday, excluding holidays, next-business-day response. Calls must be received by 3:00 PM local time in order to qualify for next-business-day response.

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

Warranty service upgrades

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

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

Maintenance service options

For additional information about IBM Power Expert Care services and support options, see announcement [LS22-0005](#), dated July 12, 2022.

Non-IBM parts service

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

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

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

Usage plan machine

No

IBM hourly service rate classification

Two

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

General terms and conditions

Field-installable features

Yes

Model conversions

No

Machine installation

Client setup. Clients are responsible for installation according to the instructions IBM provides with the machine.

Graduated program license charges apply

No

Licensed Machine Code

IBM Machine Code is licensed for use by a client 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 client. 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](#)

Machine using LMC Type Model 9105-22A

Access to Machine Code updates is conditioned on entitlement and license validation in accordance with IBM policy and practice. IBM may verify entitlement through client 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.

Educational allowance

Educational allowance: A reduced charge is available to qualified education clients. The educational allowance may not be added to any other discount or allowance.

The educational allowance is 5 percentage for the products in this announcement.

Prices

For all local charges, contact your IBM representative.

Annual minimum maintenance charges

Not applicable

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.

Regional availability

Argentina, Belize, Plurinational State of Bolivia, Brazil, Chile, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Haiti, Honduras, Mexico, Nicaragua, Panama, Paraguay, Peru, Uruguay, and Bolivarian Republic of Venezuela

Trademarks

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

IBM, Power, PowerVM, AIX, IBM Cloud, IBM Z, IBM Spectrum, IBM FlashSystem, IBM Research, IBM Watson, IBM Security, IBM Cloud Pak, QRadar, Resilient, i2, Guardium and MaaS360 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 and OpenShift are registered trademarks of Red Hat Inc. in the U.S. and other countries.

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 August 8, 2022)

Modifications made to the Description, Limitations, and Terms and conditions sections.