Family 9080+04 IBM Power E1080 Enterprise server

IBM United States Sales Manual
Revised: June 14, 2022

Table of contents

Product life cycle dates

Type Model	Announced	Available	Marketing Withdrawn	Service Discontinued
9080-HEX	2021-09-08	2021-09-17	-	-

Back to top

Abstract

Power E1080

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

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

Power E1080 brings AI to where your operational data resides

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

Each Power10 processor single chip module (SCM) contains two memory controllers. Four 10-core 3.65 - 3.90 GHz (max), four 12-core 3.6 - 4.15 GHz (max), or four 15-core 3.55 - 4.00 GHz (max) are used in each system node, providing 40 cores to a 160-core system (#EDP2), 48 cores to a 196-core system (#EDP3), or 60 cores to a 240-core system (#EDP4)(1). As few as 16 cores in the system can be activated or up to 100% of the cores in the system can be activated. Increments of one core at a time is available through built-in Capacity Upgrade on Demand (CUoD) functions to the full capacity of the system.

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

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

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

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

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

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

IBM Power Private Cloud Solution with Dynamic Capacity

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

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

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

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

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

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

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

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

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

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

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

IBM offers a Private Cloud Capacity Assessment and Implementation Service performed by Systems Lab Services

professionals, which can be preselected at time of purchase or requested for qualifying Power E1080 servers.

Power to Cloud services

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

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

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

CMC for Power

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

Power E1080 server Power10 hardware components

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

Note: (1) EDP4 is not available to order in China.

Model abstract 9080-HEX

The Power E1080 server provides the underlying Power10 hardware components:

- The most powerful and scalable server in the IBM Power portfolio:
  - Up to 240 Power 10 technology-based processor cores
  - Up to 64 TB memory
  - Up to 32 PCIe Gen4 x16 / PCIe Gen5x8 slots in system nodes
  - Up to 192 PCIe Gen3 slots with expansion drawers
  - Up to over 4,000 directly attached SAS disks or solid-state drives (SSDs)

    Up to 2,000 VMs (LPARs) per system

- System control unit, providing redundant system master Flexible Service Processor (FSP)

The Power E1080 supports:

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

Back to top

Highlights

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

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

The Power E1080 supports:

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

Back to top

Description

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

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

Power E1080feature summary

The following features are available on the Power E1080 server:

- One to four 5U system nodes
  - The Power E1080 server will support three and four systems nodes by December 10, 2021.
- One 2U system control unit
- One to four processor features per system with four single-chip modules (SCMs) per feature:

  - 3.65 - 3.90 GHz, 40-core Power10 processor (#EDP2)
  - 3.6 - 4.15 GHz, 48-core Power10 processor (#EDP3)
  - 3.55 - 4.00 GHz, 60-core Power10 processor (#EDP4)(1)
- CoD processor core activation features available on a per-core basis
- 64 DDIMM slots per system node
- DDR4 DDIMM memory cards:
  - 128 GB (4 x 32 GB), (#EMC1)
  - 256 GB (4 x 64 GB), (#EMC2)
  - 512 GB (4 x 128 GB), (#EMC3)
  - 1024 GB (4 x 256 GB), (#EMC4)
- CoD memory activation features include:
  - 100 GB Mobile Memory Activations (#EDAB)
  - 500 GB Mobile Memory Activations (#EMBK)
- AME optimized onto the processor chip (#EM8F)
- Six PCIe Gen4 x16 or PCIe Gen5 x8 and two PCIe Gen5 x8 I/O low- profile expansion slots per system node (maximum 32 in a 4-node system)
- One USB port to support external attached DVD when a USB adapter is installed in a node
- Redundant hot-swap AC power supplies in each system node drawer
- Two HMC ports in the system control unit
- Optional PCIe I/O expansion drawer with PCIe slots:
  - Zero to four drawers per system node drawer (#EMX0).
  - Each I/O drawer holds one or two 6-slot PCIe fan-out modules (#EMXH).
  - Each fan-out module attaches to the system node through a PCIe optical cable adapter (#EJ24).

System nodes

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

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

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

Processor core activations

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

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

60-core (#EDP4)(2)	48-core (#EDP3)	40-core (#EDP2)

1-core static activation (#EDPD)(2)	1-core static activation (#EDPC)	1-core static activation (#EDPB)
1-core Power Linux (#ELCM)(2)	1-core Power Linux (#ELCQ)	1-core Power Linux (#ELCL)

Note: (2) Features EDP4, EDPD, and ELCM are not available to order in China.

Memory

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

Each system node has 64 memory DDIMM slots, and at least half of the memory slots are always physically filled. Sixteen DDIMM slots are local to each of the four SCMs in the server, but SCMs and their cores have access to all the other memory in the server. When filling the other memory slots in the SCM, a quantity of four DDIMMs must be used. Thus, the DDIMM slots of the SCMs are from 50% to 100% filled. The system node (four SCMs) DDIMM slots can have 32, 36, 40, 44, 48, 52, 56, 60 and 64 DDIMMs physically installed (quad plugging rules).

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

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

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

Memory activation features are:

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

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

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

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

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

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

Power Enterprise Pools with Mobile and Shared Utility Capacity

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

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

- Capacity Credits for Power (5819-CRD) may be applied to a Power Enterprise Pool (2.0) containing a combination of Power E1080 and E980 systems. Metered Capacity resources consumed will be debited at the same rate for both E1080 and E980 systems.
- For each Base Activation feature purchased new on a Power E1080 server that is replacing a Power E980 server in the same Power Enterprise Pool, up to three (3) Base Activation features may be exchanged from the E980 server for three (3) new, corresponding Power E1080 Base Activation features, at no additional charge, when the Power E980 system is being removed from the pool.
- To support customers migrating to Power E1080 from Power E980 systems in a Power Enterprise Pool (1.0) with Mobile Capacity, a one- time migration is being enabled to allow a quantity of Power9 Mobile Processor and Memory activation features, purchased initially on a Power E980 system, to migrate and convert to similar features on a Power E1080 system, within the same Power Enterprise Pool, at no additional charge, when the Power E980 system is being removed from the pool.
- All offers of exchange of Base Capacity features and migration of Mobile Capacity features are designed to be executed via the Entitled Systems Support portal, and are subject to its availability within a country.

System control unit

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

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

System node PCIe slots

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

PCIe I/O expansion drawer

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

PCIe Gen1, Gen2, Gen3, and Gen4 and full-high adapter cards are supported. The set of full-high PCIe adapters that are supported is found in the Sales Manual, identified by feature number. See the PCI Adapter Placement manual for the 9080-HEX for details and rules associated with specific adapters supported and their supported placement in x8 or x16 slots.

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

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

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

PCIe I/O drawer attachment and cabling

- A PCIe x16 to optical CXP converter adapter (#EJ24) and 2.0 M (#ECCR), 10.0 M (#ECCY), or 20 M (#ECCZ) CXP 12X Active Optical Cables (AOC) connect the system node to a PCIe fan-out module in the I/O expansion drawer. One ECCR, ECCY, or ECCZ feature ships two AOC cables from IBM.
- The two AOC cables connect to two CXP ports on the fan-out module and to two CXP ports on the feature EJ24 adapter. The top port of the fan-out module must be cabled to the top port of the feature EJ24 port. Likewise, the bottom two ports must be cabled together.
- It is recommended but not required that one I/O drawer be attached to two different system nodes in the same server (one drawer module attached to one system node and the other drawer module attached to a different system node). This can help provide cabling for higher availability configurations.
- It is generally recommended that any attached PCIe I/O expansion drawer be located in the same rack as the Power10 server for ease of service, but expansion drawers can be installed in separate racks if the application or other rack content requires it. If you are attaching a large number of cables, such as SAS cables or CAT5/CAT6 Ethernet cables, to a PCIe I/O drawer, then it is generally better to place that feature EMX0 drawer in a separate rack for better cable management.
  Limitation: When this cable is ordered with a system in a rack specifying IBM Plant integration, IBM Manufacturing will ship SAS cables longer than three meters in a separate box and not attempt to place the cable in the rack. This is because the longer SAS cable is probably used to attach to an EXP24S drawer in a different rack.

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

EXP24SX disk/SSD drawer

- Direct attached storage is supported with the EXP24SX SFF Gen2-bay drawer (#ESLS), an expansion drawer with 24 2.5-inch form-factor SAS bays.
- The Power E1080 server supports up to 4,032 drives with a maximum of 168 EXP24SX drawers. The maximum of 16 EXP24SX drawers per PCIe I/O drawer due to cabling considerations remains unchanged.
- The EXP24SX SFF Gen2-bay drawer (#ESLS) is an expansion drawer with 24 2.5-inch form-factor SAS bays. Slot filler panels are included for empty bays when initially shipped. The EXP24SX supports up to 24 hot-swap SFF-2 SAS HDDs or SSDs. It uses only two EIA of space in a 19-inch rack. The EXP24SX includes redundant AC power supplies and uses two power cords.
- With AIX, Linux, and VIOS, you can order the EXP24SX with four sets of 6 bays, two sets of 12 bays, or one set of 24 bays (mode 4, 2, or 1). With IBM i, you can order the EXP24SX as one set of 24 bays (mode 1). Mode setting is done by IBM Manufacturing, and there is no option provided to change the mode after it is shipped from IBM. Note that when changing modes, a skilled, technically qualified person should follow the special documented procedures. Improperly changing modes can potentially destroy existing RAID sets, prevent access to existing data, or allow other partitions to access another partition's existing data. Hire an expert to assist if you are not familiar with this type of reconfiguration work.
- The EXP24SX SAS ports are attached to a SAS PCIe adapter or pair of adapters using SAS YO or X cables.

- To maximize configuration flexibility and space utilization, the system node does not have integrated SAS bays or integrated SAS controllers. PCIe adapters and the EXP24SX can be used to provide direct access storage.
- To further reduce possible single points of failure, EXP24SX configuration rules consistent with previous Power servers are used. IBM i configurations require the drives to be protected (RAID or mirroring). Protecting the drives is highly recommended, but not required for other operating systems. All Power operating system environments that are using SAS adapters with write cache require the cache to be protected by using pairs of adapters.
- It is recommended for SAS cabling ease that the EXP24SX drawer be located in the same rack in which the PCIe adapter is located. However, it is often a good availability practice to split a SAS adapter pair across two PCIe drawers/nodes for availability and that may make the SAS cabling ease recommendation difficult or impossible to implement.
- HDDs and SSDs that were previously located in POWER8 system units or in feature 5802 or 5803 12X-attached I/O drawers (SFF-1 bays) can be retrayed and placed in EXP24S drawers. See feature conversions previously announced on the POWER8 servers. Ordering a conversion ships an SFF-2 tray or carriage onto which you can place your existing drive after removing it from the existing SFF-1 tray/carriage. The order also changes the feature number so that IBM configuration tools can better interpret what is required.
- A maximum of 16 EXP24SX drawers is needed per PCIe drawer (#EMX0) to enable SAS cables to be properly handled by the feature EMX0 cable management bracket.
- A maximum of 999 SSDs can be ordered within a single order (initial or MES) of the system.

DVD and boot devices

A device capable of reading a DVD may be attached to the system and available to perform operating system installation, maintenance, problem determination, and service actions such as maintaining system firmware and I/O microcode at their latest levels. In addition, the system must be attached to a network with software such as AIX NIM server or Linux Install Manager configured to perform these functions:

1. Disk or SSD located in an EXP24S drawer attached to a PCIe adapter
2. A network through LAN adapters
3. A SAN attached to a Fibre Channel (FC) or FC over Ethernet adapters and indicated to the server by the 0837 specify feature

- Assuming option 1 above, the minimum system configuration requires at least one SAS disk drive in the system for AIX and Linux and two for IBM i. If you are using option 3 above, a disk or SSD drive is not required.
- For IBM i, a DVD drive must be available on the server when required.
- A DVD can optionally be in the system control unit, or one or more DVDs can be located in an external enclosure such as a 7226-1U3 multimedia drawer.

Racks

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

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

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

Front door options supported with Power E1080 system nodes for the 42U slim enterprise rack (7695-S42), front acoustic door #ECRA, high- end appearance front door #ECRF, cost-effective plain front door #ECRM.

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

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

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

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

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

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

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

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

A lighter, lower-cost lift tool is FC EB3Z(3) (lift tool) and EB4Z(3) (angled shelf kit for lift tool). The EB3Z(3) lift tool provides a hand crank to lift and position a server up to 400 lb. Note that a single system node can weigh up to 86.2 kg (190 lb).
Note: (3) Features EB3Z and EB4Z are not available to order in Albania, Bahrain, Bulgaria, Croatia, Egypt, Greece, Jordan, Kuwait, Kosovo, Montenegro, Morocco, Oman, UAE, Qatar, Saudi Arabia, Serbia, Slovakia, Slovenia, Taiwan, and Ukraine.

Reliability, Availability, and Serviceability

PCIe I/O expansion drawer and racks

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

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

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

Power distribution units (PDUs)

- The Power E1080 server factory that is integrated into an IBM rack uses horizontal PDUs located in the EIA drawer space of the rack instead of the typical vertical PDUs found in the side pockets of a rack. This is done to aid cable routing. Each horizontal PDU occupies 1U. Vertically mounting the PDUs to save rack space can cause cable routing challenges and interfere with optimal service access.
- When mounting the horizontal PDUs, it is a good practice to place them almost at the top or almost at the bottom of the rack, leaving 2U or more of space at the very top or very bottom open for cable management. Mounting a horizontal PDU in the middle of the rack is generally not optimal for cable management.
- Two possible PDU ratings are supported: 60A (orderable in most countries) and 30A.
  - The 60A PDU supports four system node power supplies and one I/O expansion drawer or eight I/O expansion drawers.
  - The 30A PDU supports two system node power supplies and one I/O expansion drawer or four I/O expansion drawer
- Rack-integrated system orders require at least two of either feature 7109, 7188, or 7196.

- Intelligent PDU with universal UTG0247 connector (#7109) is for an intelligent AC power distribution unit (PDU+) that enables users to monitor the amount of power being used by the devices that are plugged in to this PDU+. This AC power distribution unit provides 12 C13 power outlets. It receives power through a UTG0247 connector. It can be used for many different countries and applications by varying the PDU-to-wall power cord, which must be ordered separately. Each PDU requires one PDU-to-wall power cord. Supported power cords include the following features: 6489, 6491, 6492, 6653, 6654, 6655, 6656, 6657, 6658, or 6667.
- The PDU (#7188) mounts in a 19-inch rack and provides 12 C13 power outlets. The PDU has six 16A circuit breakers, with two power outlets per circuit breaker. System units and expansion units must use a power cord with a C14 plug to connect to the feature 7188. One of the following line cords must be used to distribute power from a wall outlet to the feature 7188: feature 6489, 6491, 6492, 6653, 6654, 6655, 6656, 6657, 6658, or 6667.
- The three-phase PDU (#7196) provides six C19 power outlets and is rated up to 48A. It has a 4.3 m (14 ft) fixed power cord to attach to the power source (IEC309 60A plug (3P+G)). A separate to-the-wall power cord is not required or orderable. Use the power cord 2.8 m (9.2 ft), drawer to wall/IBM PDU, (250V/10A) (#6665) to connect devices to this PDU. These power cords are different than the ones used on the feature 7188 and 7109 PDUs. Supported countries for the feature 7196 PDU are Antigua and Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guam, Guatemala, Haiti, Honduras, Indonesia, Jamaica, Japan, Mexico, Netherlands Antilles, Nicaragua, Panama, Peru, Puerto Rico, Surinam, Taiwan, Trinidad and Tobago, United States, and Venezuela.

System node power

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

System control unit power

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

Concurrent maintenance or hot-plug options

The following options are maintenance or hot-plug capable:

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

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

You can access hot-plug procedures in the product documentation at IBM Documentation website.

Active Memory Expansion

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

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

The Power E1080 includes a hardware accelerator designed to boost AME efficiency and uses less Power core resource. You have a great deal of control over AME usage. Each individual AIX partition can turn on or turn off AME. Control parameters set the amount of expansion desired in each partition to help control the amount of CPU used by the AME function. An IPL is required for the specific partition that is turning memory expansion. When turned on, monitoring capabilities are available in standard AIX performance tools, such as lparstat, vmstat, topas, and svmon.

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

AME is enabled by chargeable hardware feature EM89, which can be ordered with the initial order of the system or as an MES order. A software key is provided when the enablement feature is ordered, which is applied to the system node. An IPL is not required to enable the system node. The key is specific to an individual system and is permanent. It cannot be moved to a different server.

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

Active Memory Mirroring

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

IBM i operating system

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

Feature	Description
EDP2	Processor (3.65 - 3.9 GHz 40-core node) - QPRCFEAT value for the system is EDP2.
EDP3	Processor (3.6 - 4.15 GHz 48-core node) - QPRCFEAT value for the system is EDP3.
EDP4(2)	Processor (3.55 - 4.0 GHz 60-core node) - QPRCFEAT value for the system is EDP4.

Power E1080 is IBM i software tier P30

If the 5250 Enterprise Enablement function is to be used on the server, order one or more feature ED2Z or order the full system 5250 enablement feature ED30. Feature ED2Z provides one processor core's worth of 5250 capacity, which can be spread across multiple physical processor cores or multiple partitions.
Note: (2) Features EDP4, EDPD, and ELCM are not available to order in China.

Capacity Backup for IBM i

The Capacity Backup (CBU) designation can help meet your requirements for a second system to use for backup, high availability, and disaster recovery. It enables you to temporarily transfer IBM i processor license entitlements and 5250 Enterprise Enablement entitlements purchased for a primary machine to a secondary CBU-designated system.

Temporarily transferring these resources instead of purchasing them for your secondary system may result in significant savings. Processor activations cannot be transferred as part of this CBU offering, however programs such as Power Enterprise Pools are available for moving or sharing processor activations.

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

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

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

- 9080-HEX
- 9080-M9S

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

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

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

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

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

Reliability

The reliability of systems starts with components, devices, and subsystems that are designed to be highly reliable. During the design and development process, subsystems go through rigorous verification and integration testing processes. During system manufacturing, systems go through a thorough testing process ensure product quality.

Power E1080 system RAS

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

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

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

Memory subsystem RAS

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

Key features

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

Mutual surveillance

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

Environmental monitoring functions

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

Power10 processor functions

As in POWER9, the Power10 processor has the ability to do processor instruction retry for some transient errors and provide core-contained checkstops for certain solid faults.

Cache availability

The L3 caches in the Power10 processor in the memory buffer chip are protected with double-bit detect, single-bit correct error detection code (ECC). In addition, a threshold of correctable errors detected on cache lines can result in the data in the cache lines being purged and the cache lines removed from further operation without requiring a reboot in the PowerVM environment.

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

Special Uncorrectable Error handling

Special Uncorrectable Error (SUE) handling prevents an uncorrectable error in memory or modified cache data from immediately causing the system to terminate. Rather, the system tags the data and determines whether it will ever be used again. If the error is irrelevant, it will not force a check stop. When and if data is used, I/O adapters controlled by an I/O hub controller would freeze if data were transferred to an I/O device, otherwise, termination may be limited to the program/kernel or if the data is not owned by the hypervisor.

PCI extended error handling

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

Uncorrectable error recovery

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

Serviceability

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

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

- Design for SSR Setup, Install, and Service
- Error Detection and Fault Isolation (ED/FI)
- First Failure Data Capture (FFDC)
- Guiding Light service indicators
- Service labels and service diagrams available on the system and delivered through IBM Documentation
- Step-by-step service procedures documented in IBM Documentation or available through the HMC
- Automatic reporting of serviceable events to IBM through the Electronic Service Agent Call Home application

Service environment

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

The Power10 processor-based platforms support several service environments:

- Attachment to one or more HMCs or vHMCs is a supported option by the system with PowerVM. This is the default configuration for servers supporting logical partitions with dedicated or virtual I/O. In this case, all servers have at least one logical partition.

Service interface

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

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

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

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

First Failure Data Capture and error data analysis

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

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

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

Diagnostics

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

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

Automatic diagnostics

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

Standalone diagnostics with PowerVM

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

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

Concurrent maintenance

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

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

Service labels

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

- Location diagrams: Location diagrams are located on the system hardware, relating information regarding the placement of hardware components. Location diagrams may include location codes, drawings of physical locations, concurrent maintenance status, or other data pertinent to a repair. Location diagrams are especially useful when multiple components such as DIMMs, CPUs, processor books, fans, adapter cards, LEDs, and power supplies are installed.
- Remove/replace procedures: Service labels that contain remove/replace procedures are often found on a cover of the system or in other spots accessible to the servicer. These labels provide systematic procedures, including diagrams detailing how to remove or replace certain serviceable hardware components.

- Arrows: Numbered arrows are used to indicate the order of operation and the serviceability direction of components. Some serviceable parts such as latches, levers, and touch points need to be pulled or pushed in a certain direction and in a certain order for the mechanical mechanisms to engage or disengage. Arrows generally improve the ease of serviceability.

QR labels

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

Packaging for service

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

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

Error handling and reporting

In the event of system hardware or environmentally induced failure, the system runtime error capture capability systematically analyzes the hardware error signature to determine the cause of failure. The analysis result will be stored in system NVRAM. When the system can be successfully restarted either manually or automatically, or if the system continues to operate, the error will be reported to the operating system. Hardware and software failures are recorded in the system log. When an HMC is attached in the PowerVM environment, an ELA routine analyzes the error, forwards the event to the Service Focal Point (SFP) application running on the HMC, and notifies the system administrator that it has isolated a likely cause of the system problem. The service processor event log also records unrecoverable checkstop conditions, forwards them to the SFP application, and notifies the system administrator.

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

In the HMC-managed environment, a call home service request will be initiated from the HMC and the pertinent failure data with service parts information and part locations will be sent to an IBM service organization. Client contact information and specific system-related data such as the machine type, model, and serial number, along with error log data related to the failure, are sent to IBM Service.

Live Partition Mobility

With PowerVM Live Partition Mobility (LPM), users can migrate an AIX, IBM i or Linux VM running on one Power to another Power without disrupting services. The migration transfers the entire system environment, including processor state, memory, attached virtual devices, and connected users. It provides continuous operating system and application availability during planned outages for repair of hardware and firmware faults. The Power E1080 and other servers using Power10-technology processors with firmware level FW1010 or above supports secure LPM, whereby the VM image is encrypted and compressed prior to transfer. Secure LPM uses on-chip encryption and compression capabilities of the Power10 processor for optimal performance.

Service processor

Diagnostic monitoring of recoverable error from the processor chipset is performed on the system processor itself, while the fatal diagnostic monitoring of the processor chipset is performed by the service processor. It runs on its own boundary and does not require resources from a system processor to be operational to perform its tasks.

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

Call home

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

IBM Electronic Services

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

System configuration and inventory information collected by Electronic Service Agent also can be viewed on the secure Electronic Services web portal and used to improve problem determination and resolution between the client and the IBM support team. As part of an increased focus to provide even better service to IBM clients, Electronic Service Agent tool configuration and activation comes standard with the system. In support of this effort, a new HMC External Connectivity security whitepaper has been published, which describes data exchanges between the HMC and the IBM Service Delivery Center (SDC) and the methods and protocols for this exchange. To read the whitepaper and prepare for Electronic Service Agent installation, see the "Security" section at the IBM Electronic Service Agent website.

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

Benefits: increased uptime

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

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

For additional information, see the IBM Electronic Service Agent website.

More accurate reporting

Because system information and error logs are automatically uploaded to the IBM Support Center in conjunction with the service request, clients are not required to find and send system information, decreasing the risk of misreported or misdiagnosed errors. Once inside IBM, problem error data is run through a data knowledge management system, and knowledge articles are appended to the problem record.

Customized support

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

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

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

Back to top

Models

Model summary matrix

Model	Processor	Processor sockets	Memory	Hard Drive	CD-ROM
HEX	POWER10	16	64TB	64 EXP24SX Drawers, 136 SAS Drives	DVD-RAM

Customer setup (CSU)

No.

Devices supported

Not applicable.

Model conversions

Not available.

Feature conversions

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

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

The following conversions are available to customers:

Feature conversions for 9080-HEX adapters features:

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

Feature conversions for 9080-HEX cable features:

From FC:	To FC:	RETURN PARTS
ECC6 - 2M Optical Cable Pair for PCIe3 Expansion Drawer	ECCR - 2M Active Optical Cable Pair for PCIe3 Expansion Drawer	
ECC8 - 10M Optical Cable Pair for PCIe3 Expansion	ECCY - 10M Active Optical Cable Pair for PCIe3	

Drawer	Expansion Drawer
ECC9 - 20M Optical Cable Pair for PCIe3 Expansion Drawer	ECCZ - 20M Active Optical Cable Pair for PCIe3 Expansion Drawer

Feature conversions for 9080-HEX global resource activation

features:

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

Feature conversions for 9080-HEX memory features:

From FC:	To FC:	RETURN PARTS
EMQZ - 100 GB of #EMAZ Memory activation for HEX	EDAB - 100 GB DDR4 Mobile Memory Activation for HEX	No
EDA1 - 1 GB Memory activation from previous	EDAP - 1 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EMAZ - 1 GB Memory activation for HEX	EDAP - 1 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EDA2 - 100 GB Memory activation from previous	EDAQ - 100 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EMQZ - 100 GB of #EMAZ Memory activation for HEX	EDAQ - 100 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EMBZ - 512 GB Memory Activations for HEX	EDAR - 512 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EME7 - 512GB Memory Activation from previous	EDAR - 512 GB Base Memory activation (Pools 2.0) from Static/ Mobile	No
EDAB - 100 GB DDR4 Mobile Memory Activation for HEX	EDAU - 100 GB Base Memory activation (Pools 2.0) MES only	No
EDA7 - 100 GB DDR4 Mobile Memory Activation for HEX (From previous)	EDAV - 100 GB Base Memory Activation (Pools 2.0) from Mobile	No
EMBK - 500 GB DDR4 Mobile Memory Activation for HEX	EDAW - 500 GB Base Memory Activation (Pools 2.0) from Mobile	No
ELME - 512 GB Power Linux Memory Activations for HEX	EDAX - 512 GB Base Memory Activation Linux only - Conversion	No

Feature conversions for 9080-HEX processor features:

From FC:	To FC:	RETURN PARTS
EDPB - 1 core Processor Activation for #EDP2	EDPZ - Mobile processor activation for HEX	No
EDPC - 1 core Processor Activation for #EDP3	EDPZ - Mobile processor activation for HEX	No
EDPD - 1 core Processor Activation for #EDP4	EDPZ - Mobile processor activation for HEX	No
EDPZ - Mobile processor activation for HEX	EPDC - 1 core Base Processor Activation (Pools 2.0) for EDP2 any OS	No
EDPZ - Mobile processor activation for HEX	EPDD - 1 core Base Processor Activation (Pools 2.0) for EDP3 any OS	No
EDPZ - Mobile processor activation for HEX	EPDS - 1 core Base Processor Activation (Pools 2.0) for EDP4 any OS	No
EDA4 - 1 core Processor Activation for #EDP2 from previous	EPS0 - 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Static/ Mobile)	No
EDPB - 1 core Processor Activation for #EDP2	EPS0 - 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Static/ Mobile)	No
EDAC - 1 core Processor Activation for #EDP3 from previous	EPS1 - 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Static/ Mobile)	No
EDPC - 1 core Processor Activation for #EDP3	EPS1 - 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Static/ Mobile)	No
EDAD - 1 core Processor Activation for #EDP4 from previous	EPS2 - 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Static/ Mobile)	No
EDPD - 1 core Processor Activation for #EDP4/	EPS2 - 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Static/ Mobile)	No
ELCL - Power Linux processor activation for #EDP2	EPS5 - 1 core Base Proc Act (Pools 2.0) for #EDP2 Linux (from Static/ Mobile)	No
ELCQ - Power Linux processor activation for #EDP3	EPS6 - 1 core Base Proc Act (Pools 2.0) for #EDP3 Linux (from Static/ Mobile)	No
ELCM - Power Linux processor activation for #EDP4	EPS7 - 1 core Base Proc Act (Pools 2.0) for #EDP4 Linux (from Static/ Mobile)	No
EP2Y - 1 core Mobile Processor Activation	EPSK - 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Prev)	No
EP2Y - 1 core Mobile Processor Activation	EPSL - 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Prev)	No
EP2Y - 1 core Mobile	EPSM - 1 core Base Proc Act	No

Processor Activation	(Pools 2.0) for #EDP4 any OS (from Prev)

Feature conversions for 9080-HEX rack related features:

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

Back to top

Technical description

Physical specifications

Power E1080 model HEX:

System node:

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

System control unit:

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

PCIe Gen3 I/O expansion drawer:

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

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

Earthquake conformance

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

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

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

Operating environment

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

Limitations

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

Hardware requirements

Not applicable.

Software requirements

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

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

- AIX Version 7.3 with the 7300-00 Technology Level, 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 7200-04-01-1939, or later
- AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-05-1937, or later

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

- IBM i 7.5, or later
- IBM i 7.4 TR5, or later
- IBM i 7.3 TR11, or later

See the IBM Prerequisite website for compatibility information for hardware features and the corresponding AIX, IBM i and Linux Technology Levels.

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

- SUSE Linux Enterprise Server for SAP with SUSE Linux Enterprise Server 15, Service Pack 3, or later
- SUSE Linux Enterprise Server 12, Service Pack 5, or later
- SUSE Linux Enterprise Server 15, Service Pack 3, or later

- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8 for Power LE version 8.4, or later
- Red Hat Enterprise Linux 8 for Power LE, version 8.4, or later

If installing VIOS:

- VIOS 3.1.3.10, or later

The AIX System to Maps website should be referenced to identify the most recent information on recommended AIX levels for Power: https://www.ibm.com/support/pages/system-aix-maps

Java is supported on Power10 processor-based system. For the best use of the performance capabilities and the most recent improvements of Power10 technology, IBM recommends upgrading Java-based applications to Java 7, Java 8, or later, whenever possible. For those clients who want to run Java in AIX environments, see the AIX Download and service information website. For those clients who want to run Java in Linux environments, see the Linux Download information website.

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

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

Back to top

Publications

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

- Licenses, notices, safety, and warranty information
- Planning for the system
- Installing and configuring the system
- Troubleshooting, service, and support
- Installing, configuring, and managing consoles, terminals, and interfaces
- Installing operating systems
- Creating a virtual computing environment
- Enclosures and expansion units
- Glossary

The following information is shipped with the 9080-HEX:

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

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

To access the IBM Publications Center Portal, go to the IBM Publications Center website.

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

Back to top

Features

Features - No charge

- NONE

Features - Chargeable

Special Features - Initial Orders

- Administrative
  - (#B0PR) -SP Hard Drive/Media Retention - Power 5 years
  - (#B0VH) -SP HDR/MR POWER 3Y
  - (#ESWK) -AIX Update Access Key (UAK)
- Global Resource Activation
  - (#EP2X) -Lab Services Private Cloud Capacity Assessment
- Languages
  - (#9300) -Language Group Specify - US English
  - (#9700) -Language Group Specify - Dutch
  - (#9703) -Language Group Specify - French
  - (#9704) -Language Group Specify - German
  - (#9705) -Language Group Specify - Polish
  - (#9706) -Language Group Specify - Norwegian
  - (#9707) -Language Group Specify - Portuguese
  - (#9708) -Language Group Specify - Spanish
  - (#9711) -Language Group Specify - Italian
  - (#9712) -Language Group Specify - Canadian French
  - (#9714) -Language Group Specify - Japanese
  - (#9715) -Language Group Specify - Traditional Chinese (Taiwan)
  - (#9716) -Language Group Specify - Korean
  - (#9718) -Language Group Specify - Turkish
  - (#9719) -Language Group Specify - Hungarian
  - (#9720) -Language Group Specify - Slovakian
  - (#9721) -Language Group Specify - Russian
  - (#9722) -Language Group Specify - Simplified Chinese (PRC)
  - (#9724) -Language Group Specify - Czech
  - (#9725) -Language Group Specify - Romanian
  - (#9726) -Language Group Specify - Croatian
  - (#9727) -Language Group Specify - Slovenian
  - (#9728) -Language Group Specify - Brazilian Portuguese
  - (#9729) -Language Group Specify - Thai
- Linecords
  - (#END3) -Power Cable - Drawer to IBM PDU, 200-240V/10A for India

- (#END7) -Power Cable - Drawer to IBM PDU (250V/10A) for India

- Miscellaneous
  - (#1120) -Decline Electronic Service Agent Install Indicator
  - (#5000) -Software Preload Required
  - (#9461) -Month Indicator
  - (#9462) -Day Indicator
  - (#9463) -Hour Indicator
  - (#9464) -Minute Indicator
  - (#9465) -Qty Indicator
  - (#9466) -Countable Member Indicator
  - (#EHKV) -SAP HANA TRACKING FEATURE
- Rack Related
  - (#ER1A) -Specify Reserve 4 EIA Rack Space for PCIe3 Expansion Drawer
- Specify Codes
  - (#0205) -RISC-to-RISC Data Migration
  - (#4650) -Rack Indicator-Not Factory Integrated
  - (#4651) -Rack Indicator, Rack #1
  - (#4652) -Rack Indicator, Rack #2
  - (#4653) -Rack Indicator, Rack #3
  - (#4654) -Rack Indicator, Rack #4
  - (#4655) -Rack Indicator, Rack #5
  - (#4656) -Rack Indicator, Rack #6
  - (#4657) -Rack Indicator, Rack #7
  - (#4658) -Rack Indicator, Rack #8
  - (#4659) -Rack Indicator, Rack #9
  - (#4660) -Rack Indicator, Rack #10
  - (#4661) -Rack Indicator, Rack #11
  - (#4662) -Rack Indicator, Rack #12
  - (#4663) -Rack Indicator, Rack #13
  - (#4664) -Rack Indicator, Rack #14
  - (#4665) -Rack Indicator, Rack #15
  - (#4666) -Rack Indicator, Rack #16
  - (#8453) -Base Customer Spec Plcmnt
  - (#9169) -Order Routing Indicator-System Plant
  - (#9440) -New AIX License Core Counter
  - (#9441) -New IBM i License Core Counter
  - (#9442) -New Red Hat License Core Counter
  - (#9443) -New SUSE License Core Counter
  - (#9444) -Other AIX License Core Counter
  - (#9445) -Other Linux License Core Counter
  - (#9446) -3rd Party Linux License Core Counter
  - (#9447) -VIOS Core Counter
  - (#9449) -Other License Core Counter
  - (#ESCZ) -iSCSI SAN Load Source Specify for AIX

Special Features - Plant and/or Field Installable

Adapters

-
  - (#5260) -PCIe2 LP 4-port 1GbE Adapter
  - (#5899) -PCIe2 4-port 1GbE Adapter
  - (#EC2R) -PCIe3 LP 2-Port 10Gb NIC&ROCE SR/Cu Adapter
  - (#EC2S) -PCIe3 2-Port 10Gb NIC&ROCE SR/Cu Adapter
  - (#EC2T) -PCIe3 LP 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter
  - (#EC2U) -PCIe3 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter
  - (#EC67) -PCIe4 LP 2-port 100Gb ROCE EN LP adapter
  - (#EC6J) -PCIe2 LP 2-Port USB 3.0 Adapter
  - (#EC77) -PCIe4 LP 2-port 100Gb Crypto Connectx-6 DX QFSP56
  - (#EDFP) -Flexible service processor
  - (#EJ0J) -PCIe3 RAID SAS Adapter Quad-port 6Gb x8
  - (#EJ0M) -PCIe3 LP RAID SAS Adapter Quad-Port 6Gb x8
  - (#EJ10) -PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8
  - (#EJ11) -PCIe3 LP SAS Tape/DVD Adapter Quad-port 6Gb x8
  - (#EJ14) -PCIe3 12GB Cache RAID PLUS SAS Adapter Quad-port 6Gb x8
  - (#EJ24) -PCIe x16 to CXP Converter Card, Supports optical cables
  - (#EJ33) -PCIe3 Crypto Coprocessor BSC-Gen3 4767
  - (#EJ37) -PCIe3 Crypto Coprocessor BSC-Gen3 4769
  - (#EN0S) -PCIe2 4-Port (10Gb+1GbE) SR+RJ45 Adapter
  - (#EN0T) -PCIe2 LP 4-Port (10Gb+1GbE) SR+RJ45 Adapter
  - (#EN0W) -PCIe2 2-port 10/1GbE BaseT RJ45 Adapter
  - (#EN0X) -PCIe2 LP 2-port 10/1GbE BaseT RJ45 Adapter
  - (#EN1A) -PCIe3 32Gb 2-port Fibre Channel Adapter
  - (#EN1B) -PCIe3 LP 32Gb 2-port Fibre Channel Adapter
  - (#EN1C) -PCIe3 16Gb 4-port Fibre Channel Adapter
  - (#EN1D) -PCIe3 LP 16Gb 4-port Fibre Channel Adapter
  - (#EN1E) -PCIe3 16Gb 4-port Fibre Channel Adapter
  - (#EN1F) -PCIe3 LP 16Gb 4-port Fibre Channel Adapter
  - (#EN1G) -PCIe3 2-Port 16Gb Fibre Channel Adapter
  - (#EN1H) -PCIe3 LP 2-Port 16Gb Fibre Channel Adapter
  - (#EN1J) -PCIe4 32Gb 2-port Optical Fibre Channel Adapter
  - (#EN1K) -PCIe4 LP 32Gb 2-port Optical Fibre Channel Adapter
  - (#EN2A) -PCIe3 16Gb 2-port Fibre Channel Adapter
  - (#EN2B) -PCIe3 LP 16Gb 2-port Fibre Channel Adapter
- Administrative
  - (#EHS2) -SSD Placement Indicator - #ESLS/#ELLS
  - (#EME0) -100GB Static to Mobile Memory Auto Conversion
  - (#EPE0) -Static to Mobile Processor Auto Conversion
  - (#ESC0) -S&H - No Charge
  - (#ESC9) -S&H
  - (#EVSN) -Enable Virtual Serial Number
  - (#SVPC) -5000 Power to Cloud Reward points
    Feature SVPC not orderable in China
- Cable
  - (#3684) -SAS Cable (AE) Adapter to Enclosure, single controller/ single path 3M

    (#7802) -Ethernet Cable, 15m, Hardware Management Console to System Unit

-
- (#EB2J) -10m (30.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers
- (#EB2K) -30m (90.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers
- (#EB46) -10GbE Optical Transceiver SFP+ SR
- (#EB47) -25GbE Optical Transceiver SFP28
- (#EB48) -1GbE Base-T Transceiver RJ45
- (#EB49) -QSFP28 to SFP28 Connector
- (#EB4J) -0.5m SFP28/25GbE copper Cable
- (#EB4K) -1.0m SFP28/25GbE copper Cable
- (#EB4M) -2.0m SFP28/25GbE copper Cable
- (#EB4P) -2.0m QSFP28/100GbE copper split Cable to SFP28 4x25GbE
- (#EB57) -QSFP+ 40GbE Base-SR4 Transceiver
- (#EB59) -100GbE Optical Transceiver QSFP28
- (#EB5K) -1.0M 100GbE Copper Cable QSFP28
- (#EB5L) -1.5M 100GbE Copper Cable QSFP28
- (#EB5M) -2.0M 100GbE Copper Cable QSFP28
- (#EB5R) -3M 100GbE Optical Cable QSFP28 (AOC)
- (#EB5S) -5M 100GbE Optical Cable QSFP28 (AOC)
- (#EB5T) -10M 100GbE Optical Cable QSFP28 (AOC)
- (#EB5U) -15M 100GbE Optical Cable QSFP28 (AOC)
- (#EB5V) -20M 100GbE Optical Cable QSFP28 (AOC)
- (#EB5W) -30M 100GbE Optical Cable QSFP28 (AOC)
- (#EB5X) -50M 100GbE Optical Cable QSFP28 (AOC)
- (#EB5Y) -100M 100GbE Optical Cable QSFP28 (AOC)
- (#EBK4) -1.6M USB Cable
- (#EC6N) -E1080 Control Unit External USB Enablement Cable
- (#ECBJ) -SAS X Cable 3m - HD Narrow 6Gb 2-Adapters to Enclosure
- (#ECBK) -SAS X Cable 6m - HD Narrow 6Gb 2-Adapters to Enclosure
- (#ECBM) -SAS X Cable 15m - HD Narrow 3Gb 2-Adapters to Enclosure
- (#ECBT) -SAS YO Cable 1.5m - HD Narrow 6Gb Adapter to Enclosure
- (#ECBU) -SAS YO Cable 3m - HD Narrow 6Gb Adapter to Enclosure
- (#ECBV) -SAS YO Cable 6m - HD Narrow 6Gb Adapter to Enclosure
- (#ECBW) -SAS YO Cable 10m - HD Narrow 6Gb Adapter to Enclosure
- (#ECBY) -SAS AE1 Cable 4m - HD Narrow 6Gb Adapter to Enclosure
- (#ECBZ) -SAS YE1 Cable 3m - HD Narrow 6Gb Adapter to Enclosure
- (#ECCR) -2M Active Optical Cable Pair for PCIe3 Expansion Drawer
- (#ECCY) -10M Active Optical Cable Pair for PCIe3 Expansion Drawer
- (#ECCZ) -20M Active Optical Cable Pair for PCIe3 Expansion Drawer
- (#ECDJ) -3.0M SAS X12 Cable (Two Adapter to Enclosure)
- (#ECDK) -4.5M SAS X12 Active Optical Cable (Two Adapter to Enclosure)
- (#ECDL) -10M SAS X12 Active Optical Cable (Two Adapter to Enclosure)
- (#ECDT) -1.5M SAS YO12 Cable (Adapter to Enclosure)
- (#ECDU) -3.0M SAS YO12 Cable (Adapter to Enclosure)
- (#ECDV) -4.5M SAS YO12 Active Optical Cable (Adapter to Enclosure)
- (#ECDW) -10M SAS YO12 Active Optical Cable (Adapter to Enclosure)
- (#ECE0) -0.6M SAS AA12 Cable (Adapter to Adapter)

- (#ECE3) -3.0M SAS AA12 Cable
- (#ECE4) -4.5M SAS AA12 Active Optical Cable (Adapter to Adapter)
- (#ECW0) -Optical Wrap Plug
- (#EFCE) -System Node to System Control Unit Cable Set for Drawer 2
- (#EFCF) -System Node to System Control Unit Cable Set for Drawer 3
- (#EFCG) -System Node to System Control Unit Cable Set for Drawer 4
- (#EFCH) -System Node to System Control Unit Cable Set for Drawer 1
- (#EN01) -1m (3.3-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper
- (#EN02) -3m (9.8-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper
- (#EN03) -5m (16.4-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper

- Disk
  - (#1953) -300GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)
  - (#1964) -600GB 10k RPM SAS SFF-2 HDD for AIX/Linux
  - (#ESEU) -571GB 10K RPM SAS SFF-HDD 4K for IBM i
  - (#ESEV) -600GB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux
  - (#ESF2) -1.1TB 10K RPM SAS SFF-2 HDD 4K for IBM i
  - (#ESF3) -1.2TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux
  - (#ESFS) -1.7TB 10K RPM SAS SFF-2 HDD 4K for IBM i
  - (#ESFT) -1.8TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux
  - (#ESNL) -283GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)
  - (#ESNM) -300GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/ Linux)
  - (#ESNQ) -571GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)
  - (#ESNR) -600GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/ Linux)
- Drive
  - (#1107) -USB 500 GB Removable Disk Drive
  - (#EU01) -1TB Removable Disk Drive Cartridge
  - (#EU2T) -2TB Removable Disk Drive Cartridge (RDX)
- Global Resource Activation
  - (#EB35) -Mobile Enablement
  - (#EH35) -Mobile Enablement
  - (#EP20) -Power Enterprise Pools 2.0 Enablement
    Feature EP20 not orderable in China

- Linecords
  - (#6458) -Power Cord 4.3m (14-ft), Drawer to IBM PDU (250V/10A)
  - (#6460) -Power Cord 4.3m (14-ft), Drawer To OEM PDU (125V, 15A)
  - (#6469) -Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (250V/ 15A) U. S.
  - (#6470) -Power Cord 1.8m (6-ft), Drawer to Wall (125V/15A)
  - (#6471) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/ 10A)
  - (#6472) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/ 16A)
  - (#6473) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/ 10A)
  - (#6474) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/ 13A)
  - (#6475) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/ 16A)
  - (#6476) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/ 10A)
  - (#6477) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/ 16A)
  - (#6478) -Power Cord 2.7 M(9-foot), To Wall/OEM PDU, (250V, 16A)
  - (#6488) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (125V/ 15A or 250V/10A)
  - (#6489) -4.3m (14-Ft) 3PH/32A 380-415V Power Cord

- (#6491) -4.3m (14-Ft) 1PH/63A 200-240V Power Cord
- (#6492) -4.3m (14-Ft) 1PH/60A (48A derated) 200-240V Power Cord
- (#6493) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/ 10A)
- (#6494) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/ 10A)
- (#6496) -Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 10A)
- (#6577) -Power Cable - Drawer to IBM PDU, 200-240V/10A
- (#6651) -Power Cord 2.7M (9-foot), To Wall/OEM PDU, (125V, 15A)
- (#6653) -4.3m (14-Ft) 3PH/16A 380-415V Power Cord
- (#6654) -4.3m (14-Ft) 1PH/30A (24A derated) Power Cord
- (#6655) -4.3m (14-Ft) 1PH/30A (24A derated) WR Power Cord
- (#6656) -4.3m (14-Ft) 1PH/32A Power Cord
- (#6657) -4.3m (14-Ft) 1PH/32A Power Cord-Australia
- (#6658) -4.3m (14-Ft) 1PH/30A (24A derated) Power Cord-Korea
- (#6659) -Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 15A)
- (#6660) -Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (125V/ 15A)
- (#6665) -Power Cord 2.8m (9.2-ft), Drawer to IBM PDU, (250V/10A)
- (#6667) -4.3m (14-Ft) 3PH/32A 380-415V Power Cord-Australia
- (#6669) -Power Cord 4.3M (14-foot), Drawer to OEM PDU, (250V, 15A)
- (#6671) -Power Cord 2.7M (9-foot), Drawer to IBM PDU, 250V/10A
- (#6672) -Power Cord 2M (6.5-foot), Drawer to IBM PDU, 250V/10A
- (#6680) -Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/ 10A)
- (#ECJ5) -4.3m (14-Ft) PDU to Wall 3PH/24A 200-240V Delta-wired Power Cord
- (#ECJ6) -4.3m (14-Ft) PDU to Wall 3PH/40A 200-240V Power Cord
- (#ECJ7) -4.3m (14-Ft) PDU to Wall 3PH/48A 200-240V Delta-wired Power Cord
- (#ELC0) -PDU Access Cord 0.38m
- (#ELC1) -4.3m (14-Ft) PDU to Wall 24A 200-240V Power Cord North America
- (#ELC2) -4.3m (14-Ft) PDU to Wall 3PH/24A 415V Power Cord North America
- (#ELC5) -Power Cable - Drawer to IBM PDU (250V/10A)
- (#EQ77) -Qty 150 of #6577

- Manufacturing Instruction
  - (#0374) -HMC Factory Integration Specify
  - (#0375) -Display Factory Integration Specify
  - (#0376) -Reserve Rack Space for UPS
  - (#0377) -Reserve Rack Space for HMC
  - (#0378) -Reserve Rack Space for Display
  - (#EFCR) -Captive Rack identifier
  - (#EJRL) -Non-paired Indicator EJ14 PCIe SAS RAID+ Adapter
  - (#EJRU) -Non-paired Indicator EJ0L PCIe SAS RAID Adapter
  - (#EJW1) -Specify Mode-1 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)YO12 for EXP24SX #ESLS/ELLS
  - (#EJW2) -Specify Mode-1 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for EXP24SX #ESLS/ELLS
  - (#EJW3) -Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS
  - (#EJW4) -Specify Mode-2 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS
  - (#EJW5) -Specify Mode-4 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS
  - (#EJW6) -Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for EXP24SX #ESLS/ELLS
  - (#EJW7) -Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for EXP24SX #ESLS/ELLS
  - (#EJWA) -Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)YO12 for EXP24SX #ESLS/ELLS

- (#EJWB) -Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS
- (#EJWC) -Specify Mode-4 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS
- (#EJWD) -Specify Mode-4 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS
- (#EJWE) -Specify Mode-4 & (3)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS
- (#EJWF) -Specify Mode-1 & (2)EJ14 & (2)YO12 for EXP24SX #ESLS/ ELLS
- (#EJWG) -Specify Mode-2 & (2)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS
- (#EJWH) -Specify Mode-2 & (2)EJ14 & (1)X12 for EXP24SX #ESLS/ELLS
- (#EJWJ) -Specify Mode-2 & (4)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS
- (#EJWP) -Specify Mode-1 & (2)EJ0L & (2)YO12 for EXP24SX #ESLS/ ELLS
- (#EJWR) -Specify Mode-2 & (4)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS
- (#EJWS) -Specify Mode-2 & (2)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS
- (#EJWT) -Specify Mode-2 & (2)EJ0L& (1)X12 for EXP24SX #ESLS/ELLS

- Media Devices
  - (#EUA5) -Standalone USB DVD drive w/cable
- Memory
  - (#EDAB) -100 GB DDR4 Mobile Memory Activation for HEX
  - (#EDAG) -256 GB Base Memory Activation (Pools 2.0)
    Feature EDAG not orderable in China
  - (#EDAH) -512 GB Base Memory Activation (Pools 2.0)
    Feature EDAH not orderable in China
  - (#EDAL) -256 GB Base Memory Activation Linux only
    Feature EDAL not orderable in China
  - (#EDAM) -512 GB Base Memory Activation Linux only
    Feature EDAM not orderable in China
  - (#EDAP) -1 GB Base Memory activation (Pools 2.0) from Static
    Feature EDAP not orderable in China
  - (#EDAQ) -100 GB Base Memory activation (Pools 2.0) from Static
    Feature EDAQ not orderable in China
  - (#EDAR) -512 GB Base Memory activation (Pools 2.0) from Static
    Feature EDAR not orderable in China
  - (#EDAS) -500 GB Base Memory activation (Pools 2.0) from Static
    Feature EDAS not orderable in China
  - (#EDAT) -1 GB Base Memory activation (Pools 2.0) MES only
    Feature EDAT not orderable in China
  - (#EDAU) -100 GB Base Memory activation (Pools 2.0) MES only
    Feature EDAU not orderable in China
  - (#EDAV) -100 GB Base Memory Activation (Pools 2.0) from Mobile
    Feature EDAV not orderable in China
  - (#EDAW) -500 GB Base Memory Activation (Pools 2.0) from Mobile
    Feature EDAW not orderable in China
  - (#EDAX) -512 GB Base Memory Activation Linux only - Conversion
    Feature EDAX not orderable in China

- (#ELME) -512 GB Power Linux Memory Activations for HEX
- (#EM8F) -Active Memory expansion enablement for HEX
- (#EMAC) -512 GM Memory Activation for #EHC9
- (#EMAZ) -1 GB Memory activation for HEX
- (#EMBK) -500 GB DDR4 Mobile Memory Activation for HEX
- (#EMBZ) -512 GB Memory Activations for HEX
- (#EMC1) -128 GB (4x32GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory
- (#EMC2) -256 GB (4x64GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory
- (#EMC3) -512 GB (4x128GB) DDIMMs, 2933 MHz, 16GBIT DDR4 Memory
- (#EMC4) -1 TB (4x256GB) DDIMMs, 2933 MHz, 16GBIT DDR4 Memory
- (#EME6) -512GB Linux only Memory Activation (from Prev)
- (#EME7) -512GB Memory Activation from previous
- (#EMQZ) -100 GB of #EMAZ Memory activation for HEX

- Miscellaneous
  - (#0983) -US TAA Compliance Indicator
  - (#0984) -Product assembled in USA manufacturing plant
  - (#1140) -Custom Service Specify, Rochester Minn, USA
  - (#2145) -Primary OS - IBM i
  - (#2146) -Primary OS - AIX
  - (#2147) -Primary OS - Linux
  - (#4891) -CBU SPECIFY
  - (#ECSF) -Custom Service Specify, Montpellier, France
  - (#ECSM) -Custom Service Specify, Mexico
  - (#ECSP) -Custom Service Specify, Poughkeepsie, USA
  - (#ECSS) -Integrated Solution Packing
  - (#EJBC) -4-NVMe U.2 (7mm) Flash drive bays
  - (#ELCN) -PowerVM for Linux indicator for HEX
  - (#EU29) -Order Placed Indicator
  - (#SVBP) -BP Post-Sale Services: 1 Day
  - (#SVCS) -IBM Systems Lab Services Post-Sale Services: 1 Day
  - (#SVNN) -Other IBM Post-Sale Services: 1 Day
- Power
  - (#7109) -Intelligent PDU+, 1 EIA Unit, Universal UTG0247 Connector
  - (#7188) -Power Distribution Unit
  - (#7196) -Power Distribution Unit (US) - 1 EIA Unit, Universal, Fixed Power Cord
  - (#ECJJ) -High Function 9xC19 Single-Phase or Three-Phase Wye PDU plus
  - (#ECJL) -High Function 9xC19 PDU plus 3-Phase Delta
  - (#ECJN) -High Function 12xC13 Single-Phase or Three-Phase Wye PDU plus
  - (#ECJQ) -High Function 12xC13 PDU plus 3-Phase Delta
  - (#EMXA) -AC Power Supply Conduit for PCIe3 Expansion Drawer
  - (#EPTJ) -High Function 9xC19 PDU: Switched, Monitoring
  - (#EPTL) -High Function 9xC19 PDU 3-Phase: Switched, Monitoring
  - (#EPTN) -High Function 12xC13 PDU: Switched, Monitoring
  - (#EPTQ) -High Function 12xC13 PDU 3-Phase: Switched, Monitoring
  - (#ESLA) -Specify AC Power Supply for EXP12SX/EXP24SX Storage Enclosure
- Processor
  - (#ED2Z) -Single 5250 Enterprise Enablement

- (#ED30) -Full 5250 Enterprise Enablement
- (#EDP2) -40-core (4x10) 3.65 to 3.90 GHz (max) Power10 Processor with 5U system node drawer
- (#EDP3) -48-core (4x12) 3.60 to 4.15 GHz (max) Power10 Processor with 5U system node drawer
- (#EDP4) -60-core (4x15) 3.55 to 4.00 GHz (max) Power10 Processor with 5U system node drawer
  Feature EDP4 not orderable in China.

- (#EDPB) -1 core Processor Activation for #EDP2
- (#EDPC) -1 core Processor Activation for #EDP3
- (#EDPD) -1 core Processor Activation for #EDP4
  Feature EDPD not orderable in China.

- (#EDPZ) -Mobile processor activation for HEX
- (#EHC9) -Solution Edition for Healthcare 3.7 GHZ, 60-core Processor (CEC)
- (#ELCL) -Power Linux processor activation for #EDP2
- (#ELCM) -Power Linux processor activation for #EDP4
  Feature ELCM not orderable in China.

- (#ELCP) -1 core Processor activation for #EHC9 no cost
- (#ELCQ) -Power Linux processor activation for #EDP3
  Feature EME4 not orderable in China.

- (#EP2Y) -1 core Mobile Processor Activation
- (#EPDC) -1 core Base Processor Activation (Pools 2.0) for EDP2 any OS
  Feature EPDC not orderable in China

- (#EPDD) -1 core Base Processor Activation (Pools 2.0) for EDP3 any OS
  Feature EPDD not orderable in China

- (#EPDS) -1 core Base Processor Activation (Pools 2.0) for EDP4 any OS
  Feature EPDS not orderable in China.

- (#EPDU) -1 core Base Processor Activation (Pools 2.0) for EDP2 Linux only
  Feature EPDU not orderable in China

- (#EPDW) -1 core Base Processor Activation (Pools 2.0) for EDP3 Linux only
  Feature EPDW not orderable in China

- (#EPDX) -1 core Base Processor Activation (Pools 2.0) for EDP4 Linux only
  Feature EPDX not orderable in China.

- (#EPS0) -1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Static)
  Feature EPS0 not orderable in China

- (#EPS1) -1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Static)
  Feature EPS1 not orderable in China

- (#EPS2) -1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Static)
  Feature EPS2 not orderable in China.

- (#EPS5) -1 core Base Proc Act (Pools 2.0) for #EDP2 Linux (from Static)
  Feature EPS5 not orderable in China

- (#EPS6) -1 core Base Proc Act (Pools 2.0) for #EDP3 Linux (from Static)
  Feature EPS6 not orderable in China

- (#EPS7) -1 core Base Proc Act (Pools 2.0) for #EDP4 Linux (from Static)
  Feature EPS7 not orderable in China.
- (#EPSK) -1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Mobile Prev)
  Feature EPSK not orderable in China
- (#EPSL) -1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Mobile Prev)
  Feature EPSL not orderable in China
- (#EPSM) -1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Mobile Prev)
  Feature EPSM not orderable in China.

- Rack Related
  - (#7118) -Environmental Monitoring Probe
  - (#EB2Z) -Lift Tool
  - (#EB3Z) -Lift tool based on GenieLift GL-8 (standard)
  - (#EB4Z) -Service wedge shelf tool kit for EB3Z
  - (#EBAB) -IBM Rack-mount Drawer Bezel and Hardware
  - (#EBAC) -OEM Rack-mount Drawer Bezel and Hardware
  - (#EBAE) -System Node (5U) Drawer Indicator for Solution Edition for Healthcare
  - (#EDBK) -2 x SMP cable brackets for non-IBM Rack
  - (#EDN1) -5U System node Indicator drawer
  - (#EMX0) -PCIe Gen3 I/O Expansion Drawer
  - (#EMXH) -PCIe3 6-Slot Fanout Module for PCIe3 Expansion Drawer
  - (#EPTH) -Horizontal PDU Mounting Hardware
  - (#ER16) -Indicator, reserve 5 EIA rack space
  - (#ER21) -Field Integration of Rack and Server
  - (#ESLS) -EXP24SX SAS Storage Enclosure
- Services
  - (#0010) -One CSC Billing Unit
  - (#0011) -Ten CSC Billing Units
- Solid State Drive
  - (#EC5J) -Mainstream 800 GB SSD NVMe U.2 module
  - (#EC5K) -Mainstream 1.6 TB SSD NVMe U.2 module
  - (#EC5L) -Mainstream 3.2 TB SSD NVMe U.2 module
  - (#EC7Q) -800GB Mainstream NVMe U.2 SSD 4k for AIX/Linux
  - (#ELV9) -ETK9 Load Source Specify (387 GB SSD SFF-2)
  - (#ELVD) -ETKD Load Source Specify (775 GB SSD SFF-2)
  - (#ELVH) -ETKH Load Source Specify (1.55 TB SSD SFF-2)
  - (#ES94) -387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
  - (#ES95) -387GB Enterprise SAS 4k SFF-2 SSD for IBM i
  - (#ESB2) -387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
  - (#ESB6) -775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
  - (#ESBA) -387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
  - (#ESBB) -387GB Enterprise SAS 4k SFF-2 SSD for IBM i
  - (#ESBG) -775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
  - (#ESBH) -775GB Enterprise SAS 4k SFF-2 SSD for IBM i
  - (#ESBL) -1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
  - (#ESBM) -1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i

- (#ESGV) -387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
- (#ESGZ) -775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
- (#ESJ0) -931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJ1) -931GB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESJ2) -1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJ3) -1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESJ4) -3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJ5) -3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESJ6) -7.45TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJ7) -7.45TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESJJ) -931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJK) -931GB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESJL) -1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJM) -1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESJN) -3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJP) -3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESJQ) -7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESJR) -7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESK8) -387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ESK9) -387GB Enterprise SAS 4k SFF-2 SSD for IBM i
- (#ESKC) -775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ESKD) -775GB Enterprise SAS 4k SFF-2 SSD for IBM i
- (#ESKG) -1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ESKH) -1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i
- (#ESKK) -931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESKM) -931GB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESKP) -1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESKR) -1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESKT) -3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESKV) -3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESKX) -7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESKZ) -7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESMB) -931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESMD) -931GB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESMF) -1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESMH) -1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESMK) -3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESMS) -3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESMV) -7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
- (#ESMX) -7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i
- (#ESNA) -775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ESNB) -775GB Enterprise SAS 4k SFF-2 SSD for IBM i
- (#ESNE) -1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ESNF) -1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i
- (#ETK1) -387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
- (#ETK3) -775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

- (#ETK8) -387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ETK9) -387GB Enterprise SAS 4k SFF-2 SSD for IBM i
- (#ETKC) -775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ETKD) -775GB Enterprise SAS 4k SFF-2 SSD for IBM i
- (#ETKG) -1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
- (#ETKH) -1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i

- Specify Codes
  - (#0040) -Mirrored System Disk Level, Specify Code
  - (#0041) -Device Parity Protection-All, Specify Code
  - (#0047) -Device Parity RAID-6 All, Specify Code
  - (#0265) -AIX Partition Specify
  - (#0266) -Linux Partition Specify
  - (#0267) -IBM i Operating System Partition Specify
  - (#0296) -Specify Custom Data Protection
  - (#0308) -Mirrored Level System Specify Code
  - (#0347) -RAID Hot Spare Specify
  - (#0837) -SAN Load Source Specify
  - (#5550) -Sys Console On HMC
  - (#EB73) -IBM i 7.3 Indicator
  - (#EB74) -IBM i 7.4 Indicator
  - (#EB75) -IBM i 7.5 Indicator
  - (#EHR2) -Boot Drive / Load Source in EXP24SX Specify (in #ESLS or #ELLS)
  - (#EL9D) -ESMD Load Source Specify (931GB SSD SFF-2)
  - (#EL9H) -ESMH Load Source Specify (1.86TB SSD SFF-2)
  - (#EL9S) -ESMS Load Source Specify (3.72TB SSD SFF-2)
  - (#EL9X) -ESMX Load Source Specify (7.44TB SSD SFF-2)
  - (#ELKM) -ESKM Load Source Specify (931GB SSD SFF-2)
  - (#ELKR) -ESKR Load Source Specify (1.86TB SSD SFF-2)
  - (#ELKV) -ESKV Load Source Specify (3.72TB SSD SFF-2)
  - (#ELKZ) -ESKZ Load Source Specify (7.44TB SSD SFF-2)
  - (#ELT2) -#ESF2 Load Source Specify (1.1TB HDD SFF-2)
  - (#ELTS) -#ESFS Load Source Specify (1.7TB HDD SFF-2)
  - (#ELTU) -#ESEU Load Source Specify (571GB HDD SFF-2)
  - (#ELU9) -ESK9 Load Source Specify (387GB SSD SFF-2)
  - (#ELUD) -ESKD Load Source Specify (775GB SSD SFF-2)
  - (#ELUH) -ESKH Load Source Specify (1.55TB SSD SFF-2)
  - (#ELUK) -ESJK Load Source Specify (931GB SSD SFF-2)
  - (#ELUL) -#ESNL Load Source Specify (283GB HDD SFF-2)
  - (#ELUM) -ESJM Load Source Specify (1.86TB SSD SFF-2)
  - (#ELUP) -ESJP Load Source Specify (3.72TB SSD SFF-2)
  - (#ELUQ) -#ESNQ Load Source Specify (571GB HDD SFF-2)
  - (#ELUR) -ESJR Load Source Specify (7.44TB SSD SFF-2)
  - (#ELZ5) -ES95 Load Source Specify (387GB SSD SFF-2)
  - (#ELZB) -ESNB Load Source Specify (775GB SSD SFF-2)
  - (#ELZF) -ESNF Load Source Specify (1.55TB SSD SFF-2)
  - (#ESLB) -ESBB Load Source Specify (387GB SSD SFF-2)

- (#ESLH) -ESBH Load Source Specify (775GB SSD SFF-2)
- (#ESLM) -ESBM Load Source Specify (1.55TB SSD SFF-2)
- (#EU41) -ESJ1 Load Source Specify (931GB SSD SFF-2)
- (#EU43) -ESJ3 Load Source Specify (1.86TB SSD SFF-2)
- (#EU45) -ESJ5 Load Source Specify (3.72TB SSD SFF-2)
- (#EU47) -ESJ7 Load Source Specify (7.45TB SSD SFF-2)

- Virtualization Engine
  - (#5228) -PowerVM Enterprise Edition
  - (#EPA0) -Deactivation of LPM (Live Partition Mobility)

Feature availability matrix

The following feature availability matrix for MT 9080 uses the letter "A" to indicate features that are available and orderable on the specified models. "S" indicates a feature that is supported on the new model during a model conversion; these features will work on the new model, but additional quantities of these features cannot be ordered on the new model; they can only be removed. "N" indicates that the feature is not supported on the new model and must be removed during the model conversion. As additional features are announced, supported, or withdrawn, this list will be updated. Please check with your Marketing Representative for additional information.

```
        |H| A = AVAILABLE  S = SUPPORTED
        |E| N = NOT SUPPORTED, MUST BE REMOVED
        |X|
        | |
FEAT/PN | |        DESCRIPTION
--------|-|-------------------------------------------------------------

0010    |A| One CSC Billing Unit
0011    |A| Ten CSC Billing Units
0040    |A| Mirrored System Disk Level, Specify Code
0041    |A| Device Parity Protection-All, Specify Code
0047    |A| Device Parity RAID-6 All, Specify Code
0205    |A| RISC-to-RISC Data Migration
0265    |A| AIX Partition Specify
0266    |A| Linux Partition Specify
0267    |A| IBM i Operating System Partition Specify
0296    |A| Specify Custom Data Protection
0308    |A| Mirrored Level System Specify Code
0347    |A| RAID Hot Spare Specify
0374    |A| HMC Factory Integration Specify
0375    |A| Display Factory Integration Specify
0376    |A| Reserve Rack Space for UPS
0377    |A| Reserve Rack Space for HMC
0378    |A| Reserve Rack Space for Display
0837    |A| SAN Load Source Specify
0983    |A| US TAA Compliance Indicator
0984    |A| Product assembled in USA manufacturing plant
1107    |A| USB 500 GB Removable Disk Drive
1120    |A| Decline Electronic Service Agent Install Indicator
1140    |A| Custom Service Specify, Rochester Minn, USA
1953    |S| 300GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)
1964    |S| 600GB 10k RPM SAS SFF-2 HDD for AIX/Linux
2145    |A| Primary OS - IBM i
2146    |A| Primary OS - AIX
2147    |A| Primary OS - Linux
```

```
3684    |A| SAS Cable (AE) Adapter to Enclosure, single controller/
        | |single path 3M
4650    |A| Rack Indicator- Not Factory Integrated
4651    |A| Rack Indicator, Rack #1
4652    |A| Rack Indicator, Rack #2
4653    |A| Rack Indicator, Rack #3
4654    |A| Rack Indicator, Rack #4
4655    |A| Rack Indicator, Rack #5
4656    |A| Rack Indicator, Rack #6
4657    |A| Rack Indicator, Rack #7
4658    |A| Rack Indicator, Rack #8
4659    |A| Rack Indicator, Rack #9
4660    |A| Rack Indicator, Rack #10
4661    |A| Rack Indicator, Rack #11
4662    |A| Rack Indicator, Rack #12
4663    |A| Rack Indicator, Rack #13
4664    |A| Rack Indicator, Rack #14
4665    |A| Rack Indicator, Rack #15
4666    |A| Rack Indicator, Rack #16
4891    |A| CBU SPECIFY
5000    |S| Software Preload Required
5228    |A| PowerVM Enterprise Edition
5260    |S| PCIe2 LP 4-port 1GbE Adapter
5550    |A| Sys Console On HMC
5899    |S| PCIe2 4-port 1GbE Adapter
6458    |A| Power Cord 4.3m (14-ft), Drawer to IBM PDU (250V/10A)
6460    |A| Power Cord 4.3m (14-ft), Drawer To OEM PDU (125V, 15A)
6469    |A| Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (250V/15A)
        | |U. S.
6470    |A| Power Cord 1.8m (6-ft), Drawer to Wall (125V/15A)
6471    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/10A)
6472    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/16A)
6473    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/10A)
6474    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/13A)
6475    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/16A)
6476    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)
6477    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/16A)
6478    |A| Power Cord 2.7 M(9-foot), To Wall/OEM PDU, (250V, 16A)
6488    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (125V/15A
        | |or 250V/10A )
6489    |A| 4.3m (14-Ft) 3PH/32A 380-415V Power Cord
6491    |A| 4.3m (14-Ft) 1PH/63A 200-240V Power Cord
6492    |A| 4.3m (14-Ft) 1PH/60A (48A derated) 200-240V Power Cord
6493    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)
6494    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)
6496    |A| Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 10A)
6577    |A| Power Cable - Drawer to IBM PDU, 200-240V/10A
6651    |A| Power Cord 2.7M (9-foot), To Wall/OEM PDU, (125V, 15A)
6653    |A| 4.3m (14-Ft) 3PH/16A 380-415V Power Cord
6654    |A| 4.3m (14-Ft) 1PH/30A (24A derated) Power Cord
6655    |A| 4.3m (14-Ft) 1PH/30A (24A derated) WR Power Cord
6656    |A| 4.3m (14-Ft) 1PH/32A Power Cord
6657    |A| 4.3m (14-Ft) 1PH/32A Power Cord-Australia
6658    |A| 4.3m (14-Ft) 1PH/30A (24A derated) Power Cord-Korea
```

```
6659    |A| Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 15A)
6660    |A| Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (125V/15A)
6665    |A| Power Cord 2.8m (9.2-ft), Drawer to IBM PDU, (250V/10A)
6667    |A| 4.3m (14-Ft) 3PH/32A 380-415V Power Cord-Australia
6669    |A| Power Cord 4.3M (14-foot), Drawer to OEM PDU, (250V, 15A)
6671    |A| Power Cord 2.7M (9-foot), Drawer to IBM PDU, 250V/10A
6672    |A| Power Cord 2M (6.5-foot), Drawer to IBM PDU, 250V/10A
6680    |A| Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)
7109    |S| Intelligent PDU+, 1 EIA Unit, Universal UTG0247 Connector
7118    |S| Environmental Monitoring Probe
7188    |A| Power Distribution Unit
7196    |S| Power Distribution Unit (US) - 1 EIA Unit, Universal,
        | |Fixed Power Cord
7802    |S| Ethernet Cable, 15m, Hardware Management Console to System
        | |Unit
8453    |A| Base Customer Spec Plcmnt
9169    |A| Order Routing Indicator- System Plant
9300    |A| Language Group Specify - US English
9440    |A| New AIX License Core Counter
9441    |A| New IBM i License Core Counter
9442    |A| New Red Hat License Core Counter
9443    |A| New SUSE License Core Counter
9444    |A| Other AIX License Core Counter
9445    |A| Other Linux License Core Counter
9446    |A| 3rd Party Linux License Core Counter
9447    |A| VIOS Core Counter
9449    |A| Other License Core Counter
9461    |A| Month Indicator
9462    |A| Day Indicator
9463    |A| Hour Indicator
9464    |A| Minute Indicator
9465    |A| Qty Indicator
9466    |A| Countable Member Indicator
9700    |A| Language Group Specify - Dutch
9703    |A| Language Group Specify - French
9704    |A| Language Group Specify - German
9705    |A| Language Group Specify - Polish
9706    |A| Language Group Specify - Norwegian
9707    |A| Language Group Specify - Portuguese
9708    |A| Language Group Specify - Spanish
9711    |A| Language Group Specify - Italian
9712    |A| Language Group Specify - Canadian French
9714    |A| Language Group Specify - Japanese
9715    |A| Language Group Specify - Traditional Chinese (Taiwan)
9716    |A| Language Group Specify - Korean
9718    |A| Language Group Specify - Turkish
9719    |A| Language Group Specify - Hungarian
9720    |A| Language Group Specify - Slovakian
9721    |A| Language Group Specify - Russian
9722    |A| Language Group Specify - Simplified Chinese (PRC)
9724    |A| Language Group Specify - Czech
9725    |A| Language Group Specify - Romanian
9726    |A| Language Group Specify - Croatian
9727    |A| Language Group Specify - Slovenian
```

```
9728    |A| Language Group Specify - Brazilian Portuguese
9729    |A| Language Group Specify - Thai
B0PR    |A| SP Hard Drive/Media Retention - Power 5 years
B0SR    |A| ServicePac for Basic Selection for Power 5 years
B0VC    |A| ServicePac for Basic Selection for Power 3 years
B0VH    |A| SP HDR/MR POWER 3Y
EB2J    |A| 10m (30.3-ft) - IBM MTP 12 strand cable for 40/100G
        | |transceivers
EB2K    |A| 30m (90.3-ft) - IBM MTP 12 strand cable for 40/100G
        | |transceivers
EB2Z    |A| Lift Tool
EB35    |A| Mobile Enablement
EB3Z    |A| Lift tool based on GenieLift GL-8 (standard)
EB46    |A| 10GbE Optical Transceiver SFP+ SR
EB47    |A| 25GbE Optical Transceiver SFP28
EB48    |A| 1GbE Base-T Transceiver RJ45
EB49    |A| QSFP28 to SFP28 Connector
EB4J    |A| 0.5m SFP28/25GbE copper Cable
EB4K    |A| 1.0m SFP28/25GbE copper Cable
EB4M    |A| 2.0m SFP28/25GbE copper Cable
EB4P    |A| 2.0m QSFP28/100GbE copper split Cable to SFP28 4x25GbE
EB4Z    |A| Service wedge shelf tool kit for EB3Z
Feature EB3Z not orderable in Albania, Bahrain, Bulgaria,
Croatia, Egypt, Greece, Jordan, Kuwait, Kosovo, Montenegro, Morocco,
Oman, UAE, Qatar, Saudi Arabia, Serbia, Slovakia, Slovenia, Taiwan
and Ukraine
EB57    |A| QSFP+ 40GbE Base-SR4 Transceiver
EB59    |A| 100GbE Optical Transceiver QSFP28
EB5K    |A| 1.0M 100GbE Copper Cable QSFP28
EB5L    |A| 1.5M 100GbE Copper Cable QSFP28
EB5M    |A| 2.0M 100GbE Copper Cable QSFP28
EB5R    |A| 3M 100GbE Optical Cable QSFP28 (AOC)
EB5S    |A| 5M 100GbE Optical Cable QSFP28 (AOC)
EB5T    |A| 10M 100GbE Optical Cable QSFP28 (AOC)
EB5U    |A| 15M 100GbE Optical Cable QSFP28 (AOC)
EB5V    |A| 20M 100GbE Optical Cable QSFP28 (AOC)
EB5W    |A| 30M 100GbE Optical Cable QSFP28 (AOC)
EB5X    |A| 50M 100GbE Optical Cable QSFP28 (AOC)
EB5Y    |S| 100M 100GbE Optical Cable QSFP28 (AOC)
EB73    |A| IBM i 7.3 Indicator
EB74    |A| IBM i 7.4 Indicator
EB75    |A| IBM i 7.5 Indicator
EBAB    |A| IBM Rack-mount Drawer Bezel and Hardware
EBAC    |A| OEM Rack-mount Drawer Bezel and Hardware
EBAE    |A| System Node (5U) Drawer Indicator for Solution Edition for
        | |Healthcare
EBK4    |S| 1.6M USB Cable
EC2R    |S| PCIe3 LP 2-Port 10Gb NIC&ROCE SR/Cu Adapter
EC2S    |S| PCIe3 2-Port 10Gb NIC&ROCE SR/Cu Adapter
EC2T    |A| PCIe3 LP 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter
EC2U    |A| PCIe3 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter
EC5J    |S| Mainstream 800 GB SSD NVMe U.2 module
EC5K    |S| Mainstream 1.6 TB SSD NVMe U.2 module
EC5L    |S| Mainstream 3.2 TB SSD NVMe U.2 module
```

```
EC67    |A| PCIe4 LP 2-port 100Gb ROCE EN LP adapter
EC6J    |A| PCIe2 LP 2-Port USB 3.0 Adapter
EC6N    |A| E1080 Control Unit External USB Enablement Cable
EC77    |A| PCIe4 LP 2-port 100Gb Crypto Connectx-6 DX QFSP56
ECBJ    |A| SAS X Cable 3m - HD Narrow 6Gb 2-Adapters to Enclosure
ECBK    |A| SAS X Cable 6m - HD Narrow 6Gb 2-Adapters to Enclosure
ECBM    |S| SAS X Cable 15m - HD Narrow 3Gb 2-Adapters to Enclosure
ECBT    |A| SAS YO Cable 1.5m - HD Narrow 6Gb Adapter to Enclosure
ECBU    |A| SAS YO Cable 3m - HD Narrow 6Gb Adapter to Enclosure
ECBV    |A| SAS YO Cable 6m - HD Narrow 6Gb Adapter to Enclosure
ECBW    |A| SAS YO Cable 10m - HD Narrow 6Gb Adapter to Enclosure
ECBY    |A| SAS AE1 Cable 4m - HD Narrow 6Gb Adapter to Enclosure
ECBZ    |A| SAS YE1 Cable 3m - HD Narrow 6Gb Adapter to Enclosure
ECCR    |A| 2M Active Optical Cable Pair for PCIe3 Expansion Drawer
ECCY    |A| 10M Active Optical Cable Pair for PCIe3 Expansion Drawer
ECCZ    |A| 20M Active Optical Cable Pair for PCIe3 Expansion Drawer
ECDJ    |A| 3.0M SAS X12 Cable (Two Adapter to Enclosure)
ECDK    |A| 4.5M SAS X12 Active Optical Cable (Two Adapter to
        | |Enclosure)
ECDL    |A| 10M SAS X12 Active Optical Cable (Two Adapter to Enclosure)
ECDT    |A| 1.5M SAS YO12 Cable (Adapter to Enclosure)
ECDU    |A| 3.0M SAS YO12 Cable (Adapter to Enclosure)
ECDV    |A| 4.5M SAS YO12 Active Optical Cable (Adapter to Enclosure)
ECDW    |A| 10M SAS YO12 Active Optical Cable (Adapter to Enclosure)
ECE0    |A| 0.6M SAS AA12 Cable (Adapter to Adapter)
ECE3    |A| 3.0M SAS AA12 Cable
ECE4    |A| 4.5M SAS AA12 Active Optical Cable (Adapter to Adapter)
ECJ5    |A| 4.3m (14-Ft) PDU to Wall 3PH/24A 200-240V Delta-wired
        | |Power Cord
ECJ6    |A| 4.3m (14-Ft) PDU to Wall 3PH/40A 200-240V Power Cord
ECJ7    |A| 4.3m (14-Ft) PDU to Wall 3PH/48A 200-240V Delta-wired
        | |Power Cord
ECJJ    |A| High Function 9xC19 Single-Phase or Three-Phase Wye PDU
        | |plus
ECJL    |A| High Function 9xC19 PDU plus 3-Phase Delta
ECJN    |A| High Function 12xC13 Single-Phase or Three-Phase Wye PDU
        | |plus
ECJQ    |A| High Function 12xC13 PDU plus 3-Phase Delta
ECSF    |S| Custom Service Specify, Montpellier, France
ECSM    |A| Custom Service Specify, Mexico
ECSP    |A| Custom Service Specify, Poughkeepsie, USA
ECSS    |S| Integrated Solution Packing
ECW0    |A| Optical Wrap Plug
ED2Z    |A| Single 5250 Enterprise Enablement
ED30    |A| Full 5250 Enterprise Enablement
EDAB    |A| 100 GB DDR4 Mobile Memory Activation for HEX
EDAG    |A| 256 GB Base Memory Activation (Pools 2.0)
Feature EDAG not orderable in China
EDAH    |A| 512 GB Base Memory Activation (Pools 2.0)
Feature EDAH not orderable in China
EDAL    |A| 256 GB Base Memory Activation Linux only
Feature EDAL not orderable in China
EDAM    |A| 512 GB Base Memory Activation Linux only
Feature EDAM not orderable in China
```

```
EDAP    |A| 1 GB Base Memory activation (Pools 2.0) from Static
Feature EDAP not orderable in China
EDAQ    |A| 100 GB Base Memory activation (Pools 2.0) from Static
Feature EDAQ not orderable in China
EDAR    |A| 512 GB Base Memory activation (Pools 2.0) from Static
Feature EDAR not orderable in China
EDAS    |A| 500 GB Base Memory activation (Pools 2.0) from Static
Feature EDAS not orderable in China
EDAT    |A| 1 GB Base Memory activation (Pools 2.0) MES only
Feature EDAT not orderable in China
EDAU    |A| 100 GB Base Memory activation (Pools 2.0) MES only
Feature EDAU not orderable in China
EDAV    |A| 100 GB Base Memory Activation (Pools 2.0) from Mobile
Feature EDAV not orderable in China
EDAW    |A| 500 GB Base Memory Activation (Pools 2.0) from Mobile
Feature EDAW not orderable in China
EDAX    |A| 512 GB Base Memory Activation Linux only - Conversion
Feature EDAX not orderable in China
EDBK    |A| 2 x SMP cable brackets for non-IBM Rack
EDFP    |A| Flexible service processor
EDN1    |A| 5U System node Indicator drawer
EDP2    |A| 40-core (4x10) 3.65 to 3.90 GHz (max) Power10 Processor
        | |with 5U system node drawer
EDP3    |A| 48-core (4x12) 3.60 to 4.15 GHz (max) Power10 Processor
        | |with 5U system node drawer
EDP4    |A| 60-core (4x15) 3.55 to 4.00 GHz (max) Power10 Processor
        | |with 5U system node drawer
EDPB    |A| 1 core Processor Activation for #EDP2
EDPC    |A| 1 core Processor Activation for #EDP3
EDPD    |A| 1 core Processor Activation for #EDP4
Feature EDPD not orderable in China
EDPZ    |A| Mobile processor activation for HEX
EFCE    |A| System Node to System Control Unit Cable Set for Drawer 2
EFCF    |A| System Node to System Control Unit Cable Set for Drawer 3
EFCG    |A| System Node to System Control Unit Cable Set for Drawer 4
EFCH    |A| System Node to System Control Unit Cable Set for Drawer 1
EFCR    |A| Captive Rack identifier
EHC9    |A| Solution Edition for Healthcare 3.7 GHZ, 60-core Processor
        | |(CEC)
EH35    |A| Mobile Enablement
EHKV    |A| SAP HANA TRACKING FEATURE
EHR2    |A| Boot Drive / Load Source in EXP24SX Specify (in #ESLS or
        | |#ELLS)
EHS2    |A| SSD Placement Indicator - #ESLS/#ELLS
EJ0J    |A| PCIe3 RAID SAS Adapter Quad-port 6Gb x8
EJ0M    |A| PCIe3 LP RAID SAS Adapter Quad-Port 6Gb x8
EJ10    |A| PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8
EJ11    |A| PCIe3 LP SAS Tape/DVD Adapter Quad-port 6Gb x8
EJ14    |A| PCIe3 12GB Cache RAID PLUS SAS Adapter Quad-port 6Gb x8
EJ24    |A| PCIe x16 to CXP Converter Card, Supports optical cables
EJ33    |S| PCIe3 Crypto Coprocessor BSC-Gen3 4767
EJ37    |A| PCIe3 Crypto Coprocessor BSC-Gen3 4769
EJBC    |A| 4-NVMe U.2 (7mm) Flash drive bays
EJRL    |A| Non-paired Indicator EJ14 PCIe SAS RAID+ Adapter
```

```
EJRU    |A| Non-paired Indicator EJ0L PCIe SAS RAID Adapter
EJW1    |A| Specify Mode-1 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)YO12 for
        | |EXP24SX #ESLS/ELLS
EJW2    |A| Specify Mode-1 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for
        | |EXP24SX #ESLS/ELLS
EJW3    |A| Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for
        | |EXP24SX #ESLS/ELLS
EJW4    |A| Specify Mode-2 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for
        | |EXP24SX #ESLS/ELLS
EJW5    |A| Specify Mode-4 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for
        | |EXP24SX #ESLS/ELLS
EJW6    |A| Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for
        | |EXP24SX #ESLS/ELLS
EJW7    |A| Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for
        | |EXP24SX #ESLS/ELLS
EJWA    |A| Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)YO12 for
        | |EXP24SX #ESLS/ELLS
EJWB    |A| Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for
        | |EXP24SX #ESLS/ELLS
EJWC    |A| Specify Mode-4 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for
        | |EXP24SX #ESLS/ELLS
EJWD    |A| Specify Mode-4 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for
        | |EXP24SX #ESLS/ELLS
EJWE    |A| Specify Mode-4 & (3)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for
        | |EXP24SX #ESLS/ELLS
EJWF    |A| Specify Mode-1 & (2)EJ14 & (2)YO12 for EXP24SX #ESLS/ELLS
EJWG    |A| Specify Mode-2 & (2)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS
EJWH    |A| Specify Mode-2 & (2)EJ14 & (1)X12 for EXP24SX #ESLS/ELLS
EJWJ    |A| Specify Mode-2 & (4)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS
EJWP    |S| Specify Mode-1 & (2)EJ0L & (2)YO12 for EXP24SX #ESLS/ELLS
EJWR    |S| Specify Mode-2 & (4)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS
EJWS    |S| Specify Mode-2 & (2)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS
EJWT    |S| Specify Mode-2 & (2)EJ0L& (1)X12 for EXP24SX #ESLS/ELLS
EL9D    |A| ESMD Load Source Specify (931GB SSD SFF-2)
EL9H    |A| ESMH Load Source Specify (1.86TB SSD SFF-2)
EL9S    |A| ESMS Load Source Specify (3.72TB SSD SFF-2)
EL9X    |A| ESMX Load Source Specify (7.44TB SSD SFF-2)
ELC0    |A| PDU Access Cord 0.38m
ELC1    |A| 4.3m (14-Ft) PDU to Wall 24A 200-240V Power Cord North
        | |America
ELC2    |A| 4.3m (14-Ft) PDU to Wall 3PH/24A 415V Power Cord North
        | |America
ELC5    |A| Power Cable - Drawer to IBM PDU (250V/10A)
ELCL    |A| Power Linux processor activation for #EDP2
ELCM    |A| Power Linux processor activation for #EDP4
Feature ELCM not orderable in China
ELCN    |A| PowerVM for Linux indicator for HEX
ELCQ    |A| Power Linux processor activation for #EDP3
ELKM    |A| ESKM Load Source Specify (931GB SSD SFF-2)
ELKR    |A| ESKR Load Source Specify (1.86TB SSD SFF-2)
ELKV    |A| ESKV Load Source Specify (3.72TB SSD SFF-2)
ELKZ    |A| ESKZ Load Source Specify (7.44TB SSD SFF-2)
ELME    |A| 512 GB Power Linux Memory Activations for HEX
ELCP    |A| 1 core Processor activation for #EHC9 no cost
```

```
ELT2    |A| #ESF2 Load Source Specify (1.1TB HDD SFF-2)
ELTS    |A| #ESFS Load Source Specify (1.7TB HDD SFF-2)
ELTU    |A| #ESEU Load Source Specify (571GB HDD SFF-2)
ELU9    |A| ESK9 Load Source Specify (387GB SSD SFF-2)
ELUD    |A| ESKD Load Source Specify (775GB SSD SFF-2)
ELUH    |A| ESKH Load Source Specify (1.55TB SSD SFF-2)
ELUK    |S| ESJK Load Source Specify (931GB SSD SFF-2)
ELUL    |A| #ESNL Load Source Specify (283GB HDD SFF-2)
ELUM    |S| ESJM Load Source Specify (1.86TB SSD SFF-2)
ELUP    |S| ESJP Load Source Specify (3.72TB SSD SFF-2)
ELUQ    |A| #ESNQ Load Source Specify (571GB HDD SFF-2)
ELUR    |S| ESJR Load Source Specify (7.44TB SSD SFF-2)
ELV9    |A| ETK9 Load Source Specify (387 GB SSD SFF-2)
ELVD    |A| ETKD Load Source Specify (775 GB SSD SFF-2)
ELVH    |A| ETKH Load Source Specify (1.55 TB SSD SFF-2)
ELZ5    |S| ES95 Load Source Specify (387GB SSD SFF-2)
ELZB    |S| ESNB Load Source Specify (775GB SSD SFF-2)
ELZF    |S| ESNF Load Source Specify (1.55TB SSD SFF-2)
EM8F    |A| Active Memory expansion enablement for HEX
EMAC    |A| 512 GM Memory Activation for #EHC9
EMAZ    |A| 1 GB Memory activation for HEX
EMBK    |A| 500 GB DDR4 Mobile Memory Activation for HEX
EMBZ    |A| 512 GB Memory Activations for HEX
EMC1    |A| 128 GB (4x32GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory
EMC2    |A| 256 GB (4x64GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory
EMC3    |A| 512 GB (4x128GB) DDIMMs, 2933 MHz, 16GBIT DDR4 Memory
EMC4    |A| 1 TB (4x256GB) DDIMMs, 2933 MHz, 16GBIT DDR4 Memory
EME0    |A| 100GB Static to Mobile Memory Auto Conversion
EMQZ    |A| 100 GB of #EMAZ Memory activation for HEX
EMX0    |A| PCIe Gen3 I/O Expansion Drawer
EMXA    |A| AC Power Supply Conduit for PCIe3 Expansion Drawer
EMXH    |A| PCIe3 6-Slot Fanout Module for PCIe3 Expansion Drawer
EN01    |A| 1m (3.3-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper
EN02    |A| 3m (9.8-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper
EN03    |A| 5m (16.4-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper
EN0S    |S| PCIe2 4-Port (10Gb+1GbE) SR+RJ45 Adapter
EN0T    |S| PCIe2 LP 4-Port (10Gb+1GbE) SR+RJ45 Adapter
EN0W    |A| PCIe2 2-port 10/1GbE BaseT RJ45 Adapter
EC7Q    |A| 800GB Mainstream NVMe U.2 SSD 4k for AIX/Linux
EN0X    |A| PCIe2 LP 2-port 10/1GbE BaseT RJ45 Adapter
EN1A    |A| PCIe3 32Gb 2-port Fibre Channel Adapter
EN1B    |A| PCIe3 LP 32Gb 2-port Fibre Channel Adapter
EN1C    |A| PCIe3 16Gb 4-port Fibre Channel Adapter
EN1D    |A| PCIe3 LP 16Gb 4-port Fibre Channel Adapter
EN1E    |A| PCIe3 16Gb 4-port Fibre Channel Adapter
EN1F    |A| PCIe3 LP 16Gb 4-port Fibre Channel Adapter
EN1G    |A| PCIe3 2-Port 16Gb Fibre Channel Adapter
EN1H    |A| PCIe3 LP 2-Port 16Gb Fibre Channel Adapter
EN1J    |A| PCIe4 32Gb 2-port Optical Fibre Channel Adapter
EN1K    |A| PCIe4 LP 32Gb 2-port Optical Fibre Channel Adapter
EN2A    |A| PCIe3 16Gb 2-port Fibre Channel Adapter
EN2B    |A| PCIe3 LP 16Gb 2-port Fibre Channel Adapter
EP20    |A| Power Enterprise Pools 2.0 Enablement
Feature EP20 not orderable in China
```

```
EP2X    |A| Lab Services Private Cloud Capacity Assessment
EPA0    |A| Deactivation of LPM (Live Partition Mobility)
EPDC    |A| 1 core Base Processor Activation (Pools 2.0) for EDP2 any
        | |OS
Feature EPDC not orderable in China
EPDD    |A| 1 core Base Processor Activation (Pools 2.0) for EDP3 any
        | |OS
Feature EPDD not orderable in China
EPDS    |A| 1 core Base Processor Activation (Pools 2.0) for EDP4 any
Feature EPDS not orderable in China
        | |OS
EPDU    |A| 1 core Base Processor Activation (Pools 2.0) for EDP2
        | |Linux only
Feature EPDU not orderable in China
EPDW    |A| 1 core Base Processor Activation (Pools 2.0) for EDP3
        | |Linux only
Feature EPDW not orderable in China
EPDX    |A| 1 core Base Processor Activation (Pools 2.0) for EDP4
Feature EPDX not orderable in China
        | |Linux only
EPE0    |A| Static to Mobile Processor Auto Conversion
EPS0    |A| 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from
        | |Static)
Feature EPS0 not orderable in China
EPS1    |A| 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from
        | |Static)
Feature EPS1 not orderable in China
EPS2    |A| 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from
Feature EPS2 not orderable in China
        | |Static)
EPS5    |A| 1 core Base Proc Act (Pools 2.0) for #EDP2 Linux (from
        | |Static)
Feature EPS5 not orderable in China
EPS6    |A| 1 core Base Proc Act (Pools 2.0) for #EDP3 Linux (from
        | |Static)
Feature EPS6 not orderable in China
EPS7    |A| 1 core Base Proc Act (Pools 2.0) for #EDP4 Linux (from
Feature EPS7 not orderable in China
        | |Static)
EPSK    |A| 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from
        | |Mobile Prev)
Feature EPSK not orderable in China
EPSL    |A| 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from
        | |Mobile Prev)
Feature EPSL not orderable in China
EPSM    |A| 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from
        | |Mobile Prev)
Feature EPSM not orderable in China
EPTH    |A| Horizontal PDU Mounting Hardware
EPTJ    |S| High Function 9xC19 PDU: Switched, Monitoring
EPTL    |S| High Function 9xC19 PDU 3-Phase: Switched, Monitoring
EPTN    |S| High Function 12xC13 PDU: Switched, Monitoring
EPTQ    |S| High Function 12xC13 PDU 3-Phase: Switched, Monitoring
EQ77    |S| Qty 150 of #6577
```

```
ER16    |A| Indicator, reserve 5 EIA rack space
ER1A    |A| Specify Reserve 4 EIA Rack Space for PCIe3 Expansion Drawer
ER21    |A| Field Integration of Rack and Server
ES94    |S| 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ES95    |S| 387GB Enterprise SAS 4k SFF-2 SSD for IBM i
ESB2    |S| 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ESB6    |S| 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ESBA    |S| 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESBB    |S| 387GB Enterprise SAS 4k SFF-2 SSD for IBM i
ESBG    |S| 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESBH    |S| 775GB Enterprise SAS 4k SFF-2 SSD for IBM i
ESBL    |S| 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESBM    |S| 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i
ESC0    |A| S&H - No Charge
ESC9    |A| S&H
ESCZ    |A| iSCSI SAN Load Source Specify for AIX
ESEU    |A| 571GB 10K RPM SAS SFF- HDD 4K for IBM i
ESEV    |A| 600GB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux
ESF2    |A| 1.1TB 10K RPM SAS SFF-2 HDD 4K for IBM i
ESF3    |A| 1.2TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux
ESFS    |A| 1.7TB 10K RPM SAS SFF-2 HDD 4K for IBM i
ESFT    |A| 1.8TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux
ESGV    |S| 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ESGZ    |S| 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ESJ0    |S| 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJ1    |S| 931GB Mainstream SAS 4k SFF-2 SSD for IBM i
ESJ2    |S| 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJ3    |S| 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESJ4    |S| 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJ5    |S| 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESJ6    |S| 7.45TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJ7    |S| 7.45TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESJJ    |S| 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJK    |S| 931GB Mainstream SAS 4k SFF-2 SSD for IBM i
ESJL    |S| 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJM    |S| 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESJN    |S| 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJP    |S| 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESJQ    |S| 7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESJR    |S| 7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESK1    |A| 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ESK3    |A| 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ESK8    |A| 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESK9    |A| 387GB Enterprise SAS 4k SFF-2 SSD for IBM i
ESKC    |A| 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESKD    |A| 775GB Enterprise SAS 4k SFF-2 SSD for IBM i
ESKG    |A| 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESKH    |A| 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i
ESKK    |A| 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESKM    |A| 931GB Mainstream SAS 4k SFF-2 SSD for IBM i
ESKP    |A| 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESKR    |A| 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESKT    |A| 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESKV    |A| 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
```

```
ESKX    |A| 7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESKZ    |A| 7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESLA    |A| Specify AC Power Supply for EXP12SX/EXP24SX Storage
        | |Enclosure
ESLB    |S| ESBB Load Source Specify (387GB SSD SFF-2)
ESLH    |S| ESBH Load Source Specify (775GB SSD SFF-2)
ESLM    |S| ESBM Load Source Specify (1.55TB SSD SFF-2)
ESLS    |A| EXP24SX SAS Storage Enclosure
ESMB    |A| 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESMD    |A| 931GB Mainstream SAS 4k SFF-2 SSD for IBM i
ESMF    |A| 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESMH    |A| 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESMK    |A| 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESMS    |A| 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESMV    |A| 7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux
ESMX    |A| 7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i
ESNA    |S| 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESNB    |S| 775GB Enterprise SAS 4k SFF-2 SSD for IBM i
ESNE    |S| 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ESNF    |S| 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i
ESNL    |A| 283GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)
ESNM    |A| 300GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/
        | |Linux)
ESNQ    |A| 571GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)
ESNR    |A| 600GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/
        | |Linux)
ESWK    |A| AIX Update Access Key (UAK)
ETK1    |A| 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ETK3    |A| 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux
ETK8    |A| 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ETK9    |A| 387GB Enterprise SAS 4k SFF-2 SSD for IBM i
ETKC    |A| 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ETKD    |A| 775GB Enterprise SAS 4k SFF-2 SSD for IBM i
ETKG    |A| 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux
ETKH    |A| 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i
EU01    |A| 1TB Removable Disk Drive Cartridge
EU29    |A| Order Placed Indicator
EU2T    |A| 2TB Removable Disk Drive Cartridge (RDX)
EU41    |S| ESJ1 Load Source Specify (931GB SSD SFF-2)
EU43    |S| ESJ3 Load Source Specify (1.86TB SSD SFF-2)
EU45    |S| ESJ5 Load Source Specify (3.72TB SSD SFF-2)
EU47    |S| ESJ7 Load Source Specify (7.45TB SSD SFF-2)
EUA5    |A| Standalone USB DVD drive w/cable
EVSN    |A| Enable Virtual Serial Number
EXA2    |A| 2 YEAR, ADVANCED EXPERT CARE
EXA3    |A| 3 YEAR, ADVANCED EXPERT CARE
EXA4    |A| 4 YEAR, ADVANCED EXPERT CARE
EXA5    |A| 5 YEAR, ADVANCED EXPERT CARE
EXP1    |A| 1 YEAR, PREMIUM EXPERT CARE
EXP2    |A| 2 YEAR, PREMIUM EXPERT CARE
EXP3    |A| 3 YEAR, PREMIUM EXPERT CARE
EXP4    |A| 4 YEAR, PREMIUM EXPERT CARE
EXP5    |A| 5 YEAR, PREMIUM EXPERT CARE
SVBP    |A| BP Post-Sale Services: 1 Day
```

```
SVCS    |A| IBM Systems Lab Services Post-Sale Services: 1 Day
SVNN    |A| Other IBM Post-Sale Services: 1 Day
SVPC    |A| 5000 Power to Cloud Reward points
Feature SVPC not orderable in China
```

Feature descriptions

Note: Not all of the following features are available in all countries. Check with your country representative for specific feature availability.
The following is a list of all feature codes in numeric order for the IBM Power Systems 9080 machine type.

Attributes, as defined in the following feature descriptions, state the interaction of requirements among features.

Minimums and maximums are the absolute limits for a single feature without regard to interaction with other features. The maximum valid quantity for MES orders may be different than for initial orders. The maximums listed below refer to the largest quantity of these two possibilities.

The order type defines if a feature is orderable only on initial orders, only on MES orders, on both initial and MES orders, or if a feature is supported on a model due to a model conversion. Supported features cannot be ordered on the converted model, only left on or removed from the converted model.

(#0010) One CSC Billing Unit

One Billing Unit used by the Customer Solution Center.

- Attributes provided: One CSC Billing Unit
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0011) Ten CSC Billing Units

Ten Billing Units used by the Customer Solutions Center.

- Attributes provided: Ten CSC Billing Units
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0040) Mirrored System Disk Level, Specify Code

This code indicates the level of disk protection desired and helps ensure that adequate hardware is in the final configuration.

- Attributes provided: Device-level mirrored protection
- Attributes required: Minimum of two (2) disk units

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - AIX - not supported
  - IBM i supported
  - Linux - not supported li.Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#0041) Device Parity Protection-All, Specify Code

This code indicates the level of disk protection desired and helps ensure that adequate hardware is in the final configuration.

- Attributes provided: RAID Data Protection
- Attributes required: RAID-capable disk unit controller
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - AIX - not supported
  - IBM i supported
  - Linux - not supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#0047) Device Parity RAID-6 All, Specify Code

This code indicates the level of disk protection desired and helps ensure that adequate hardware is in the final configuration.

- Attributes provided: RAID-6 Data Protection
- Attributes required: RAID-6 capable disk unit controller
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - AIX - not supported
  - IBM i supported
  - Linux - not supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#0205) RISC-to-RISC Data Migration

#0205 is used on initial orders to designate that the new server should only receive a partial load of IBM i in IBM Mfg.

When #0205 is on the order, manufacturing will only load SLIC and up through QSYS of IBM i. Ensure that enough storage is ordered to contain the additional OS code that will be loaded following installation of the system at the Client

location. Specify code #0205 is mutually exclusive with #5000, SW Preload Required

The migration process requires that the installed model be at the same version and release level of IBM i and other licensed programs as the new server.

More information, and an updated IBM i Upgrade and Data Migration Road Map (RISC-RISC) are available at

http://publib.boulder.ibm.com/iseries/

- Attributes provided: Partial load of IBM i in IBM Mfg.
- Attributes required: #2145 - Primary OS - IBM i and partition specify code #0267 and RISC to RISC Data Migration from Clients existing system
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#0265) AIX Partition Specify

This feature indicates customers intend to create a partition on the system that will use the AIX operating system. This feature should be included once for each intended AIX partition. This feature is an indicator and does not deliver parts, software, or services.

- Attributes provided: None
- Attributes required: Customers intend to create a partition on the system that will run the AIX operating system.
- Minimum required: 0
- Maximum allowed: 1000 (Initial order maximum: 1000)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0266) Linux Partition Specify

This feature indicates customers intend to create a partition on the system that will use the Linux operating system. This feature should be included once for each intended Linux partition. This feature is an indicator and does not deliver parts, software, or services.

- Attributes provided: None
- Attributes required: Customers intend to create a partition on the system that will run the Linux operating system.
- Minimum required: 0
- Maximum allowed: 1000 (Initial order maximum: 1000)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0267) IBM i Operating System Partition Specify

This feature indicates customers intend to create a partition on the system that will use the IBM i operating system. This feature should be included once for each intended IBM i partition. This feature is an indicator and does not deliver parts,

software, or services.

- Attributes provided: None
- Attributes required: Customers intend to create a partition on the system that will run the IBM i operating system.
- Minimum required: 0
- Maximum allowed: 1000 (Initial order maximum: 1000)
- OS level required:
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0296) Specify Custom Data Protection

#0296 specifies that a system has multiple IBM i partitions and that data protection schemes should be considered separately for each partition instead of only for an overall system level. Each partition's data protection scheme can be different or the same.

- Attributes provided: N/A
- Attributes required: N/A
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0308) Mirrored Level System Specify Code

This specify code indicates the level of disk protection desired and helps ensure that adequate hardware is in the final configuration.

For new systems, the marketing configurator will show an error if sufficient disk units and disk controllers are not included on the order to support IOA-level mirroring protection. #0308 causes all disk units to be placed into configurations capable of IOA-level mirroring. Each disk unit and its mirrored pair must be on a different disk controller.

Note that the load source disk unit in a new, preloaded system will be device-level mirrored (same protection as provided with feature #0040). This means that the load source is controlled by the first disk unit controller on the first system bus, and will be mirrored with a like disk unit that is also attached to the same first disk controller on the first system bus.

For upgrade orders, #0308 will cause the marketing configurator to show an error if sufficient disk units and disk controllers are not available to provide the capability to enable IOA-level mirrored protection for all DASD.

It is the client's responsibility to start mirroring on their system.

- Attributes provided: IOA level system mirroring
- Attributes required: A minimum of two disk controllers and an even number of disk units (with a minimum of four disk units on a system).
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0347) RAID Hot Spare Specify

#0347 is a specify code that indicates to IBM configuration tools and to IBM Manufacturing that RAID-5 or RAID-6 disk arrays should be further protected using the IBM i function of RAID hot spare. If specified, IBM will ship a configuration which has at least one stand- by disk drive for each disk controller in the system or designated partition. The customer may alter the hot spare configuration selecting different options once the system is installed.

- Attributes provided: N/A
- Attributes required: Existence of #0041 or #0047
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0374) HMC Factory Integration Specify

Used in manufacturing to ensure that an HMC ordered from IBM under a separate machine type/model is associated with the system order and is shipped concurrently.

- Attributes provided: N/A
- Attributes required: Ordered machine type-model HMC.
- Minimum required: 0
- Maximum allowed: 42 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: Does not apply
  Note: Available only when ordered with 19 inch feature code rack on an MES order. Max. is 36 with rack feature 0551.

(#0375) Display Factory Integration Specify

Used in manufacturing to ensure that a display ordered from IBM under a separate machine type/model, is associated with the system order, and is shipped concurrently.

- Attributes provided: N/A
- Attributes required: Ordered machine type-model display.
- Minimum required: 0
- Maximum allowed: 42 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: Does not apply
  Note: Available only when ordered with 19 inch feature code rack on an MES order. Max. is 36 with rack feature 0551.

(#0376) Reserve Rack Space for UPS

Used in manufacturing to reserve 1 EIA of rack space in the bottom of the rack for later client installation of a UPS.

- Attributes provided: 1 EIA rack space reserved.
- Attributes required: Ordered rack feature.
- Minimum required: 0
- Maximum allowed: 42 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: Does not apply
  Note: Available only when ordered with 19 inch feature code rack on an MES order. Max. is 36 with rack feature 0551.

(#0377) Reserve Rack Space for HMC

Used in manufacturing to reserve 1 EIA of rack space in the middle of the rack for later client installation of a rack-mounted HMC.

- Attributes provided: 1 EIA rack space reserved.
- Attributes required: Ordered rack feature.
- Minimum required: 0
- Maximum allowed: 42 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: Does not apply
  Note: Available only when ordered with 19 inch feature code rack on an MES order. Max. is 36 with rack feature 0551.

(#0378) Reserve Rack Space for Display

Used in manufacturing to reserve 1 EIA of rack space in the middle of the rack for later client installation of an HMC rack-mounted display such as the 7316.

- Attributes provided: 1 EIA rack space reserved.
- Attributes required: Ordered rack feature.
- Minimum required: 0
- Maximum allowed: 42 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: Does not apply

(#0837) - Fiber Channel SAN Load Source Specify

Indicates that a SAN drive is being used as the Load Source for the operating system.

- Attributes provided: SAN load source placement specify
- Attributes required: Fiber Channel adapter
- Minimum required: 0
  Maximum allowed: 1 (Initial order maximum: 1)

-
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#0983) US TAA Compliance Indicator

This feature indicates that the product was assembled in a manufacturing plant in the USA or in a country approved under the US Trade Agreement Act. Only valid on U.S. orders.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#0984) Product assembled in USA manufacturing plant

This feature indicates that the product was assembled in a manufacturing plant in the USA.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#1107) USB 500 GB Removable Disk Drive

Provides an RDX disk drive in a rugged cartridge to be used in an RDX Internal and External docking station such as the #1103, #1104, #1123, #EU03, #EU04, #EU23 or #EU07. 500 GB is uncompressed. With typical 2X compression, capacity would be 1000 GB. Compression/ decompression is provided by the operating system, not the drive itself. Feature 1107 is not entitled under the IBM Maintenance Agreement, if one is purchased.

- Attributes provided: 500 GB RDX rugged disk/cartridge
- Attributes required: None.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: See RDX Docking Station
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#1120) Decline Electronic Service Agent Install Indicator

Electronic Service Agent is a basic component of the Power Systems server and is an important tool in maintaining the

high availability expected from your IBM system and the service response time you expect from your IBM service team. You should select this feature only if you do not intend to have the IBM SSR install ESA at the time your system is being installed. Before selecting this feature please discuss this decision with your IBM Account Team or IBM Business Partner.

Electronic Service Agent installation by an IBM SSR is available to you for no additional charge during initial server installation only.

- Attributes provided: Documents Customers intention to not allow the IBM SSR to install Electronic Service Agent during system installation
- Attributes required: none
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#1140) Custom Service Specify, Rochester Minn, USA

Having #1140 on the order, will cause the order to be routed to Rochester and the machine to be internally routed to the CSC build area in building 114 (Rochester).

- Attributes provided: Customization
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: N/A
- Return parts MES: Does not apply

(#1953) 300GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)

300 GB SFF 15k RPM SAS disk drive mounted in a Gen-2 carrier and supported in SFF SAS bay such as used in the #5887 EXP24S I/O drawer or #ESLS EXP24SX I/O drawer. Disk can be formatted for either 512 bytes or 528 byte sectors and is thus called having 5xx drive sectors. 528 byte sectors provide additional protection. IBM Manufacturing will ship pre-formatted with 528 or with 512 byte sectors. Selection of the formatting is selected by IBM Manufacturing based on manufacturing rules, but the client may change at their location. In the EPX24SX IBM Manufacturing will ship using 528 byte sectors. Capacity is 300GB with 512 byte formatting and is 283GB with 528 byte sector. CCIN is 19B1.

Limitations: physical difference in carriers prevent this drive from being used in SFF-1 bays such as used in the #5802/5803 I/O drawer or in SFF-3 bays such as used in POWER8/POWER9 system units.

- Attributes provided: 300GB/283GB of SFF (2.5-inch) SAS disk storage mounted in Gen-2 carrier.
- Attributes required: one SFF-2 drive bay.
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 0)
- OS level required:
  - Linux - supported
  - AIX - supported
  - IBM i not supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes

  Return parts MES: No

- supported in #5887

(#1964) 600GB 10k RPM SAS SFF-2 HDD for AIX/Linux

600 GB SFF 10k RPM SAS disk drive mounted in a Gen-2 carrier and supported in SFF SAS bay such as used in the #5887 EXP24S I/O drawer or #ESLS EXP24SX I/O drawer. Disk can be formatted for either 512 bytes or 528 byte sectors and is thus called having 5xx drive sectors. 528 byte sectors provide additional protection. IBM Manufacturing will ship pre-formatted with 528 or with 512 byte sectors. Selection of the formatting is selected by IBM Manufacturing based on manufacturing rules, but the client may change at their location. In the EPX24SX IBM Manufacturing will ship using 528 byte sectors. Capacity is 600GB with 512 byte formatting and is 571GB with 528 byte sector.

Features #1962 and #1964 are physically identical drives with the same 19B3 CCIN. Different feature codes help the IBM configuration tools understand how the HDD is used. #1964 indicates usage by AIX, Linux or VIOS. #1962 indicates usage by IBM i.

Limitation: physical difference in carriers prevent this drive from being used in SFF-1 bays such as used in the #5802/5803 I/O drawer or in SFF-3 bays such as used in POWER8/POWER9 system units.

- Attributes provided: 600GB/571GB of SFF (2.5-inch) SAS disk storage mounted in Gen-2 carrier.
- Attributes required: one SFF-2 drive bay.
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 0)
- OS level required:
  - Linux - supported
  - AIX - supported
  - IBM i not supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#2145) Primary OS - IBM i

Indicates clients intend to use the IBM i operating system on the primary system partition. This feature is used as a Manufacturing Routing indicator and does not deliver parts, software or services.

- Attributes provided: None
- Attributes required: Indicates clients intend to use the IBM i operating system on the primary system partition.
- Minimum required: 0
- Maximum allowed: 99 (Initial order maximum: 99)
- OS level required:
  - Linux not supported
  - AIX not supported
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Only one of features 2145, 2146 or 2147 may exist on a system Configuration report.

(#2146) Primary OS - AIX

Indicates clients intend to use the AIX operating system on the primary system partition. This feature is used as a Manufacturing Routing indicator and does not deliver parts, software or services.

- Attributes provided: None
- Attributes required: Indicates clients intend to use the AIX operating system on the primary system partition.
- Minimum required: 0
- Maximum allowed: 99 (Initial order maximum: 99)
- OS level required:
  - Linux - not supported
  - AIX - supported
  - IBM i not supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Only one of features 2145, 2146 or 2147 may exist on a system Configuration report.

(#2147) Primary OS - Linux

Indicates clients intend to use the Linux operating system on the primary system partition. This feature is used as a Manufacturing Routing indicator and does not deliver parts, software or services.

- Attributes provided: None
- Attributes required: Indicates clients intend to use the Linux operating system on the primary system partition.
- Minimum required: 0
- Maximum allowed: 99 (Initial order maximum: 99)
- OS level required:
  - Linux supported
  - AIX not supported
  - IBM i not supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Only one of features 2145, 2146 or 2147 may exist on a system Configuration report.

(#3684) SAS Cable (AE) Adapter to Enclosure, single controller/single path 3M

This adapter-to-enclosure (AE) SAS cable most commonly connects a SAS controller to a media expansion drawer.

For AIX and Linux, this cable can also be used to connect two SAS adapters to a SAS disk drawer in a specific dual controller HA two system JBOD configuration using two #5912 controllers. Single controller/single path connections are supported with this cable only for this specific JBOD configuration, and, as such, two #5912 SAS controllers and two (AE style) cables are required for a supported configuration. The two SAS adapters must be in different host systems/partitions.

This cable has one mini SAS 4X plug connector on the adapter end wired in 4x mode and one mini SAS 4X plug connector on the drawer end, wired in 4x mode. Follow the directions on the connector labels when attaching the connectors on this cable. This cable is three meters long. Select the SAS (AE) cable length that best matches the distance between the host system and the remote SAS drawer being attached.

- Attributes provided: Connection between a SAS controller and a media expansion drawer or for AIX and Linux connection between #5912 SAS controller and a SAS disk drawer in a dual controller HA two system JBOD configuration only
- Attributes required: none
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)

- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#4650) Rack Indicator- Not Factory Integrated

This indicator is used to specify that the rack mountable device in this initial order should not be merged into a rack within IBM Manufacturing. If a device with 4650 is ordered with a rack, the device will not be factory integrated in the ordered rack and will ship uninstalled in the rack.
Note: This "no additional charge" feature will be placed on an initial order for a rack mountable device by the Configuration Tool when the order does not ship from IBM Manufacturing in a Rack. This server is not designed to be rack shippable.

- Attributes provided: System will not be shipped in a rack.
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply
  Note: One and only one rack indicator feature is required on all orders (#4650 to #4666).

(#4651) Rack Indicator, Rack #1

When added to an initial rack order, this indicator is used to specify the first rack for a multi rack order, or the only rack for a single rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #1.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack Integration/ Rack Specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4652) Rack Indicator, Rack #2

When added to an initial rack order, this indicator is used to specify the second rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #2 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack Integration/Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4653) Rack Indicator, Rack #3

When added to an initial rack order, this indicator is used to specify the third rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #3 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4654) Rack Indicator, Rack #4

When added to an initial rack order, this indicator is used to specify the fourth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #4 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4655) Rack Indicator, Rack #5

When added to an initial rack order, this indicator is used to specify the fifth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #5 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4656) Rack Indicator, Rack #6

When added to an initial rack order, this indicator is used to specify the sixth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #6 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4657) Rack Indicator, Rack #7

When added to an initial rack order, this indicator is used to specify the seventh rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #7 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
  CSU: N/A

-
- Return parts MES: Does not apply

(#4658) Rack Indicator, Rack #8

When added to an initial rack order, this indicator is used to specify the eighth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #8 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4659) Rack Indicator, Rack #9

When added to an initial rack order, this indicator is used to specify the ninth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #9 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4660) Rack Indicator, Rack #10

When added to an initial rack order, this indicator is used to specify the tenth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #10 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify

-
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4661) Rack Indicator, Rack #11

When added to an initial rack order, this indicator is used to specify the eleventh rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #11 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4662) Rack Indicator, Rack #12

When added to an initial rack order, this indicator is used to specify the twelfth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #12 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4663) Rack Indicator, Rack #13

When added to an initial rack order, this indicator is used to specify the thirteenth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #13 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4664) Rack Indicator, Rack #14

When added to an initial rack order, this indicator is used to specify the fourteenth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #14 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4665) Rack Indicator, Rack #15

When added to an initial rack order, this indicator is used to specify the fifteenth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #15 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
  CSU: N/A

-
- Return parts MES: Does not apply

(#4666) Rack Indicator, Rack #16

When added to an initial rack order, this indicator is used to specify the sixteenth rack for a multi rack order.

When added to an initial rack mountable device order, this indicator is used to specify that the rack mountable device (such as a system or I/O drawer) is to be mounted in rack #16 of a multi rack order.
Note: For 19" rack mountable device orders: One feature code from the group 4650 to 4666 must be listed on the order. More than one feature code from this group is not allowed. For 19" rack orders: If IBM Mfg. is to assemble a rack mountable device into the rack, one feature code selection from the group 4651 to 4666 must be listed on the order. More than one feature code selection from this group is not allowed. The quantity of this selected feature code on the 19" rack order must equal the number of rack mountable devices to be installed in the rack by IBM Mfg.

- Attributes provided: Rack specify
- Attributes required: Rack
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#4891) CBU SPECIFY

This specify code indicates this system has been properly registered as a Capacity BackUp system and has, through that registration been authorized to temporarily receive IBM i Operating System License Entitlements and 5250 Processor Enablement entitlements, from a primary system under the conditions specified at the time the system was registered. This feature is an indicator only, authorization to use this system as a backup is obtained only by registering the system with IBM on the CBU website at: www.ibm.com/systems/power/ hardware/cbu

- Attributes provided: Indicates the system has been registered for use as a CBU system for IBM i License entitlement purposes.
- Attributes required: # 2145 Primary OS - IBM i
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - Linux not supported
  - AIX not supported
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#5000) Software Preload Required

Indicates that preloaded software and/or consolidated I/O is shipped with the initial order. A maximum of one (#5000) is supported. This feature has country-specific usage.
Note: Not supported in Brazil and India.

- Attributes provided: Software Pre-load
- Attributes required: N/A
- Minimum required: 0

- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - Linux not supported
  - AIX supported
  - IBM i supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: No

(#5228) PowerVM Enterprise Edition

This feature allows the customer to create partitions that are in units of less than 1 CPU (sub-CPU LPARs) and allows the same system I/O to be virtually allocated to these partitions. When PowerVM is installed in the system, all activated processors must have the PowerVM feature. A fully activated 4-core system requires that four of this feature be ordered. An encrypted key is supplied to the customer and is installed on the system, authorizing the partitioning at the sub-processor level. PowerVM Enterprise Edition also includes Live Partition Mobility, which allows for the movement of a logical partition from one Power8, Power9 or Power10 with no application downtime.
Note: If feature 5228 is ordered, the quantity ordered must be equal to the number of active processors.

- Attributes provided: Capability to partition processor
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#5260) PCIe2 LP 4-port 1GbE Adapter

This short PCIe Gen2 adapter provides four 1Gb Ethernet ports that can be configured to run at 1000, 100 or 10 Mbps. 4-pair CAT-5 Unshielded Twisted Pair (UTP) cables up to 100 meters in length are attached to the copper RJ45 connectors. Each port is independent of one another and supports full-duplex or half-duplex. 1000 Mbps speed is not supported in Half Duplex (HDX) mode.

Feature #5260 and #5899 are electronically identical and have the same CCIN of 576F. #5260 indicates a low profile tail stock while #5899 indicates a full high tail stock.

Details for the ports include:

- AIX NIM support
- IEEE 802.3ab (1 GbE), 802.1p priority, 802.1Q VLAN tagging, 802.3x flow control, 802.3ad load-balancing and failover,
- Link aggregation, IEEE 802.3ad 802.3
- Multiple MAC addresses per interface
- MSI-X, MSI and support of legacy pin interrupts
- Ether II and IEEE 802.3 encapsulated frames
- Jumbo frames up to 9.6 Kbytes

  TCP checksum offload for IPv4 and IPv6

-
- TCP segmentation Offload (TSO) for IPv4 and IPv6
- UDP checksum offload for IPv4 and IPv6
- AIX, IBM i and Linux provide software iSCSI support through the adapter. Linux can also leverage adapter hardware support including initiator and header & data digest (CRC) generation and checking
- Attributes provided: Four-port 1 Gb Ethernet
- Attributes required: 1 Low Profile (LP) PCIe slot (Gen1 or Gen2)
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 0)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No
  Note: VIOS supported

(#5550) Sys Console On HMC

With #5550, system console function is driven by the Hardware Management Console (HMC) connected to the system. The HMC is required if the following functions are desired/selected for the system:

- Attributes provided: System Console on Hardware Management Console(HMC)
- Attributes required: Hardware Management Console (HMC)
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - IBM i - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#5899) PCIe2 4-port 1GbE Adapter

This short PCIe Gen2 adapter provides four 1Gb Ethernet ports that can be configured to run at 1000, 100 or 10 Mbps. 4-pair CAT-5 Unshielded Twisted Pair (UTP) cables up to 100 meters in length are attached to the copper RJ45 connectors. Each port is independent of one another and supports full-duplex or half-duplex. 1000 Mbps speed is not supported in Half Duplex (HDX) mode.

Feature #5260 and #5899 are electronically identical and have the same CCIN of 576F. #5260 indicates a low profile tail stock while #5899 indicates a full high tail stock.

Details for the ports include: for 5260 & 5899

- AIX NIM support
- IEEE 802.3ab (1 GbE), 802.1p priority, 802.1Q VLAN tagging, 802.3x flow control, 802.3ad load-balancing and failover,
- Link aggregation, IEEE 802.3ad 802.3
- Multiple MAC addresses per interface

- MSI-X, MSI and support of legacy pin interrupts
- Ether II and IEEE 802.3 encapsulated frames
- Jumbo frames up to 9.6 Kbytes
- TCP checksum offload for IPv4 and IPv6
- TCP segmentation Offload (TSO) for IPv4 and IPv6
- UDP checksum offload for IPv4 and IPv6
- AIX, IBM i and Linux provide software iSCSI support through the adapter. Linux can also leverage adapter hardware support including initiator and header & data digest (CRC) generation and checking
- Attributes provided: Four-port 1 Gb Ethernet
- Attributes required: 1 Full High Profile PCIe slot (Gen1 or Gen2)
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 0)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No
  - VIOS supported

(#6458) Power Cord 4.3m (14-ft), Drawer to IBM PDU (250V/10A)

Standard IBM rack power cable that goes from the system or I/O drawer to the rack power distribution unit (PDU). Cable has C13 on one end (for C14 power supply connector on system unit or I/O drawer) and C14 on the other end (for IBM PDU C13 receptacle). Note for different length C13/C14 cables see #6671 (2.7M) or #6672 (2.0M).

- Attributes provided: Power jumper cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6460) Power Cord 4.3m (14-ft), Drawer To OEM PDU (125V, 15A)

This power cord goes from the system or I/O drawer to the rack OEM power distribution unit or wall socket outlet. Cable has C13 on one end (for C14 power supply connector on system unit or I/O drawer) and plug type #4 (NEMA 5-15) on the other end.

The following countries/regions use the #6460 power cord to power the system and/or peripheral features requiring a power cord: United States, Antigua & Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, Bolivia, Bonaire, Calicos Islands, Canada, Cayman Islands, Colombia, Costa Rica, Cuba, Curacao, Dominican Republic, Ecuador, El Salvador, Guam, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Mexico, Micronesia, Montserrat, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, St. Kitts/Nevis, St. Martin, Taiwan, Tortola (BVI), Trinidad/Tobago, Venezuela.

- Attributes provided: Power cord.
- Attributes required: None.
- Minimum required: 0

- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6469) Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (250V/15A) U.S.

This power cord goes from the system or I/O drawer to the wall or rack OEM power distribution unit. Cable has C13 on one end (for C14 power supply connector on system unit or I/O drawer) and plug type #5 (NEMA 6-15) on the other end for wall or OEM PDU.

The following countries/regions use the #6469 power cord to power the system and/or peripheral features requiring a power cord:

United States, Anguilla, Antigua & Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, Bolivia, Bonaire, Caicos Is., Canada, Cayman Islands, Colombia, Costa Rica, Cuba, Curacao, Dominican Republic, Ecuador, El Salvador, Guam, Guatemala, Haiti, Honduras, Jamaica, Japan, Micronesia, Montserrat, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, St. Marten NA, Taiwan, Tortola (BVI), Thailand, Venezuela.

- Attributes provided: Power cord.
- Attributes required: None.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6470) Power Cord 1.8m (6-ft), Drawer to Wall (125V/15A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #4 (NEMA 5-15). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types. 6-foot length.

The following countries/regions use the #6470 power cord to power the system and/or peripheral features requiring a power cord:

United States, Antigua and Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, Bolivia, Bonaire, Calicos Islands, Canada, Cayman Islands, Colombia, Costa Rica, Cuba, Curacao, Dominican Republic, Ecuador, El Salvador, Guam, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Mexico, Micronesia, Montserrat, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, St. Kitts/Nevis, St. Martin, Taiwan, Tortola (BVI), Trinidad/Tobago, Venezuela.

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6471) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/10A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #73 (InMetro NBR 14136). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6471 power cord to power the system and/or peripheral features requiring a power cord:

Brazil

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6472) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/16A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #18 (CEE 7 VII). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6472 power cord to power the system and/or peripheral features requiring a power cord:

Afghanistan, Albania, Algeria, Andorra, Angola, Armenia, Austria, Belarus, Belgium, Benin, Bosnia/Herzegovina, Bulgaria, Burkina Faso, Burundi, Cambodia, Cameroon, Cape Verde, Central African Republic, Chad, Comoros, Congo, Croatia, Czech Republic, Dahomey, Djibouti, Egypt, Equatorial Guinea, Eritrea, Estonia, Ethiopia, Finland, France, French Polynesia, French Guyana, Gabon, Georgia, Germany, Greece, Guadeloupe, Guinea, Guinea-Bissau, Hungary, Iceland, Indonesia, Iran, Ivory Coast, Kazakhstan, Kyrgyzstan, Laos, Latvia, Lebanon, Lithuania, Luxembourg, Macau, Macedonia, Mali, Martinique, Mauritania, Mauritius, Mayotte, Moldova, Monaco, Mongolia, Morocco, Mozambique, Netherlands, New Caledonia, Niger, North Korea (C19 only), Norway, Poland, Portugal, Principe, Reunion, Romania, Russia, Rwanda, St. Thomas, Saudi Arabia, Senegal, Serbia, Slovenia, Somalia, South Korea (C19 only), Spain, Surinam, Sweden, Syria, Tahiti, Tajikistan, Togo, Tunisia, Turkey, Turkmenistan, Ukraine, Upper Volta, Uzbekistan, Vanuatu, Vietnam, Wallis & Futuna, Zaire, Zimbabwe.

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6473) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU (250V/10A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #19 (CEE). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6473 power cord to power the system and/or peripheral features requiring a power cord:

Denmark

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6474) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/13A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #23 (BS 1364A). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6474 power cord to power the system and/or peripheral features requiring a power cord:

Abu Dhabi, Bahrain, Botswana, Brunei, Channel Islands, Cyprus, Dominica, Gambia, Grenada, Grenadines, Guyana, Hong Kong, Iraq, Ireland, Jordan, Kenya, Kuwait, Liberia, Malawi, Malaysia, Malta, Myanmar, Nigeria, Oman, Qatar, Sierra Leone, Singapore, St. Kitts, St. Lucia, Seychelles, Sudan, Tanzania, Trinidad & Tobago, United Arab Emirates, United Kingdom, Yemen, Zambia

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6475) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/16A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #32 (SII 32-1971). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6475 power cord to power the system and/or peripheral features requiring a power cord:

Israel

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6476) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #24 (SEV 24507). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6476 power cord to power the system and/or peripheral features requiring a power cord:

Lichtenstein, Switzerland

- Attributes provided: Power cord

- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6477) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/16A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #22 (SABS 164). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6477 power cord to power the system and/or peripheral features requiring a power cord:

Bangladesh, LeSotho, Maceo, Maldives, Nambia, Pakistan, Samoa, South Africa, Sri Lanka, Swaziland, Uganda.

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6478) Power Cord 2.7 M(9-foot), To Wall/OEM PDU, (250V, 16A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #25 (CEI 23-16). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types. The following countries/regions use the #6478 power cord to power the system and/or peripheral features requiring a power cord: Chile Italy Libya

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6488) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (125V/15A or 250V/10A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. 125V, 15A or 250V, 10A, Plug Type #2. Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6488 power cord to power the system and/or peripheral features requiring a power cord:

Argentina, Paraguay, Uruguay.

- Attributes provided: Power cord
- Attributes required: None

- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6489) 4.3m (14-Ft) 3PH/32A 380-415V Power Cord

#6489 is a 14-FT/4.3m 3PH/32A power cable with a Type 46 plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6491) 4.3m (14-Ft) 1PH/63A 200-240V Power Cord

#6491 is a 14-FT/4.3m 200-240V/63A power cord with a Type 46 plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6492) 4.3m (14-Ft) 1PH/60A (48A derated) 200-240V Power Cord

Feature #6492 is a 14-FT/4.3m 200-240V/48-60A power cord with a Type 46 plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: Power Cord PDU to wall
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6493) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #62 (GB 1053). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6493 power cord to power the system and/or peripheral features requiring a power cord:

People's Republic of China.

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6494) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #69 (IS 6538). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6494 power cord to power the system and/or peripheral features requiring a power cord:

India

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6496) Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 10A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #66 (KETI). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types. The following countries/regions use the #6496 power cord to power the system and/or peripheral features requiring a power cord: North Korea South Korea

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6577) Power Cable - Drawer to IBM PDU, 200-240V/10A

This feature permits manufacturing to select the optimum PDU power jumper cord length (1.0M, 2.0M, 2.7M, or 4.3M) for rack integration. This feature is mandatory on initial order specifying factory integration with IBM racks (such as with

7965-S42 racks). Feature is not valid on initial order with non-factory integrated feature 4650. Power jumper cord has C13 on one end (for C14 power supply connector on system unit or I/O drawer) and C14 on the other end (for C13 PDU receptacle).
Note: This feature is not used for MES orders except for bulk orders by SDI clients only. See C13/C14 jumper cord features #6458 (4.3M), #6671 (2.7M), #6672 (2.0M) when not using factory integration.

- Attributes provided: One power jumper cord.
- Attributes required: At least one rack and the absence of #4650.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: 1 or 2 per I/O drawer or CEC. MES orderable for SDI's only. The MES order will ship the 14 foot cable equivalent to feature number 6458.

(#6651) Power Cord 2.7M (9-foot), To Wall/OEM PDU, (125V, 15A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #75 (KETI). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types. The following countries/regions use the #6651 power cord to power the system and/or peripheral features requiring a power cord: Taiwan

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6653) 4.3m (14-Ft) 3PH/16A 380-415V Power Cord

#6653 is a 14-FT/4.3m 3PH/16A power cord with a Type 46 plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6654) 4.3m (14-Ft) 1PH/30A (24A derated) Power Cord

Feature #6654 is a 14-FT/4.3m 200-240V/24A-30A locking power cord with a Type 12 plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: power cord
- Attributes required: None
- Minimum required: 0

- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6655) 4.3m (14-Ft) 1PH/30A (24A derated) WR Power Cord

Feature #6655 is a 14-FT/4.3m 200-240V/24A-30A water-resistant power cord with a Type 40 plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: Power Cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6656) 4.3m (14-Ft) 1PH/32A Power Cord

#6656 is a 14-FT/4.3m 200-240V/32A power cord with a Type 46 plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: PDU power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6657) 4.3m (14-Ft) 1PH/32A Power Cord-Australia

This power cord provides power to a #5889, #7188 #9188, #7109, #EPTG, #EPTM, #EPTJ, #ECJM, #ECJG, #ECJJ, #ECJN, or #EPTN power distribution unit. It connects to a wall power outlet with a PDL plug.

- Attributes provided: Power connection for a PDU
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6658) 4.3m (14-Ft) 1PH/30A (24A derated) Power Cord-Korea

This power cord provides power to a #5889, #7188, #9188, #7109, #EPTG, #EPTM, #EPTJ, #ECJM, #ECJG, #ECJJ, #ECJN, or #EPTN power distribution unit. It connects to a wall power outlet with a Korean plug.

- Attributes provided: PDU power cable
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6659) Power Cord 2.7M (9-foot), To Wall/OEM PDU, (250V, 15A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #76 (KETI). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types. The following countries/regions use the #6659 power cord to power the system and/or peripheral features requiring a power cord: Taiwan

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6660) Power Cord 4.3m (14-ft), Drawer to Wall/OEM PDU (125V/15A)

This power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #59 (NEMA 5-15). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types. 14-foot length.

This power cord meets the DENAN marking requirement in Japan.

- Attributes provided: Power Cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6665) Power Cord 2.8m (9.2-ft), Drawer to IBM PDU, (250V/10A)

Standard IBM rack power jumper cord that goes from the system or I/O drawer to the rack power distribution unit (PDU). Cable has C13 on one end (for C14 power supply connector on system unit or I/O drawer) and C20 on the other end (for IBM PDU C19 receptacle).
Note: For power jumper cord which attach to PDUs with C13 receptacles, use features such as #6577, #6458, #6671, or #6672.

- Attributes provided: Power jumper cord.
- Attributes required: None.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)

  OS level required: None

-
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6667) 4.3m (14-Ft) 3PH/32A 380-415V Power Cord-Australia

#6667 is a 14-FT/4.3m 380-45V/32A power cord with a Type PDL plug which distributes power from a power source to a Power Distribution Unit.

- Attributes provided: PDU power cable
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6669) Power Cord 4.3M (14-foot), Drawer to OEM PDU, (250V, 15A)

This power cord goes from the system or I/O drawer to the rack power distribution unit. Plug type #57 (NEMA 6-15). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types. 14-foot length. This power cord meets the DENAN marking requirement in Japan.

- Attributes provided: Power Cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6671) Power Cord 2.7M (9-foot), Drawer to IBM PDU, 250V/10A

Standard IBM rack power cable that goes from the system or I/O drawer to the rack power distribution unit (PDU). Cable has C13 on one end (for C14 power supply connector on system unit or I/O drawer) and C14 on the other end (for IBM PDU C13 receptacle). Note for different length C13/C14 cables see #6458 (4.3M) or #6672 (2.0M).

- Attributes provided: Power jumper cord.
- Attributes required: None.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6672) Power Cord 2M (6.5-foot), Drawer to IBM PDU, 250V/10A

Standard rack power cable that goes from the system or I/O drawer to the rack power distribution unit (PDU). Cable has

C13 on one end (for C14 power supply connector on system unit or I/O drawer) and C14 on the other end (for IBM PDU C13 receptacle). Note for different length C13/C14 cables see #6458 (4.3M) or #6671 (2.7M).

- Attributes provided: Power jumper cord.
- Attributes required: None.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#6680) Power Cord 2.7m (9-ft), Drawer to Wall/OEM PDU, (250V/10A)

This insulated power cord goes from the system and/or peripheral features to a wall-type outlet. Plug type #6 (AS 3112-1964 NZS 198). Refer to Corporate Bulletin C-B-2-4700-009 for a description of plug types.

The following countries/regions use the #6680 power cord to power the system and/or peripheral features requiring a power cord:

Australia, Fiji Islands, Kiribati, Nauru, New Zealand, Papua New Guinea, W. Samoa.

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#7109) Intelligent PDU+, 1 EIA Unit, Universal UTG0247 Connector

(No longer available as of April 24, 2020)

This feature is for an intelligent AC power distribution unit (PDU+) that will allow the user to monitor the amount of power being used by the devices that are plugged in to this PDU+. This AC power distribution unit provides twelve C13 power outlets. It receives power through a UTG0247 connector. It can be used for many different countries and applications by varying the PDU to Wall Power Cord, which must be ordered separately. Each PDU requires one PDU to Wall Power Cord. Supported power cords include the following features: #6489, #6491, #6492, #6653, #6654, #6655, #6656, #6657, #6658, #ELC1 or #ELC2.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Twelve C13 outlets with Power Monitoring Capability
- Attributes required: none
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No
  Note: When purchased on an MES order with a feature code rack. This PDU will be mounted in the rear side pockets

until all 4 side pockets on the rack have been filled. Any additional PDUs on the order will be mounted in 1 unit of EIA rack space.
Note: When purchased as an MES order for addition to a rack in the field. This PDU may not fit in the side pockets of your rack due to a hardware interference with the rack, and may require mounting in 1 unit of rack EIA space. Insure rack space is available before placing the MES order for this PDU when it is being ordered for field installation.

(#7118) Environmental Monitoring Probe

(No Longer Available as of December 14, 2021)

The Environmental Monitoring Probe (EMP) enables you to remotely monitor environmental conditions. Using a standard Web browser, you can view the ambient temperature and humidity of the remote environment, as well as the status of two additional contact devices, such as a smoke detector or open-door sensor. The temperature/humidity probe plugs into a RJ45 connector an a PDU+. The EMP can be used with any Powerware UPS equipped with a 10/100 Mb ConnectUPS Web/SNMP Card (firmware 3.01 or higher). The EMP can be located up to 20m (65.6 feet) away.

- Attributes provided: Monitoring of temperature, humidity, and status of two contacts/ sensors. A one meter cat5 Ethernet cable, double sided hook and loop fabric, often called VELCRO(R) tape, two tie-wraps, and screw with wall anchor for mounting.
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#7188) Power Distribution Unit

An AC Power Distribution Unit (PDU) which mounts in a 19" rack and provides twelve C13 power outlets. The #7188 has six 16A circuit breakers, with two power outlets per circuit breaker. System units and/or expansion units must use a power cord with a C14 plug to connect to the #7188.

One of the following line cords must be used to distribute power from a wall outlet to the #7188; #6489, #6491, #6492, #6653, #6654, #6655, #6656, #6657, #6658, #ELC1 or #ELC2.

- #6489 - 14-Ft 3PH/32A Power Cord
- #6491 - 14-Ft 1PH/63A Power Cord
- #6492 - 14-Ft 1PH/48-60A Power Cord
- #6653 - 14-Ft 3PH/16A Power Cord
- #6654 - 14-Ft 1PH/24-30A Power Cord
- #6655 - 14-Ft 1PH/24-30A WR Power Cord
- #6656 - 14-Ft 1PH/32A Power Cord
- #6657 - 14-Ft 1PH/24A Power Cord
- #6658 - 14-Ft 1PH/24A Power Cord-Korea
- Attributes provided: Power Distribution Unit with Twelve C13 power outlets.
- Attributes required: none
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#7196) Power Distribution Unit (US) - 1 EIA Unit, Universal, Fixed Power Cord

This AC power distribution unit provides six C19 power outlets. Fixed power cord (IEC309 60A plug (3P+G). This PDU requires 3-phase electrical service.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Six C19 power outlets
- Attributes required: 3 phase electical service
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#7802) Ethernet Cable, 15m, Hardware Management Console to System Unit

This feature provides a fifteen meter long Ethernet cable for attachment of a Hardware Management Console to the system unit.

- Attributes provided: 15M Ethernet Cable
- Attributes required: Ethernet port on Hardware Management Console
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#8453) Base Customer Spec Plcmnt

1. Requests that IBM deliver the system to the customer according to the slot in drawer hardware placement defined by the account team.
2. Eliminates the need to have these parts relocated in the customers environment as may happen if the order is placed without this feature code.
3. Client placement specifications are collected using the System Planning Tool (SPT) and processed through the marketing configurator. (Use of the SPT is not required).
4. Requires account team to submit the output of the marketing configurator into IBM manufacturing via the CSP Web site http://www.ibm.com/eserver/power/csp (US Business Partners and Distributors can bypass this step.)
5. Requires account team to assure that the marketing configurator output submitted reflects the actual order placed.

- Attributes provided: I/O component placement
- Attributes required: Marketing Configurator output submitted to the CSP Web site (US Business Partners and Distributors can bypass this step.)
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: NA
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9169) Order Routing Indicator- System Plant

This feature will be auto-selected by the Configurator Tool when required. Use of this feature will affect the routing of the order. Selection of this indicator will direct the order to a system plant for fulfillment.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9300) Language Group Specify - US English

English language group for nomenclature and standard publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9440) New AIX License Core Counter

This feature is used to count the number of cores licensed to run AIX.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9441) New IBM i License Core Counter

This feature is used to count the number of cores licensed to run IBM i.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A

- Return parts MES: Does not apply

(#9442) New Red Hat License Core Counter

This feature is used to count the number of cores licensed to run Red Hat Linux.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9443) New SUSE License Core Counter

This feature is used to count the number of cores licensed to run SUSE Linux.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9444) Other AIX License Core Counter

This feature is used to count the number of existing AIX licenses transferred from another server.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9445) Other Linux License Core Counter

This feature is used to count the number of existing Linux licenses transferred from another server.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial

- CSU: N/A
- Return parts MES: Does not apply

(#9446) 3rd Party Linux License Core Counter

This feature is used to count the number of cores licensed to run 3rd party Linux.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9447) VIOS Core Counter

This feature is used to count the number of cores licensed to run VIOS (Virtual I/O Server).

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9449) Other License Core Counter

This feature is used to count the number of other cores licensed.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9461) Month Indicator

This month indicator is used to create a date stamp to enable CFR splitting and rejoining in order to circumvent the AAS maximum limitation of 30 systems entered on any one order. The quantity ordered for this feature is generated by eConfig.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0

  Maximum allowed: 12 (Initial order maximum: 12)

-
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9462) Day Indicator

This day indicator is used to create a date stamp to enable CFR splitting and rejoining in order to circumvent the AAS maximum limitation of 30 systems entered on any one order. The quantity ordered for this feature is generated by eConfig.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 31 (Initial order maximum: 31)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9463) Hour Indicator

This hour indicator is used to create a time stamp to enable CFR splitting and rejoining in order to circumvent the AAS maximum limitation of 30 systems entered on any one order. The quantity ordered for this feature is generated by eConfig.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 24 (Initial order maximum: 24)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9464) Minute Indicator

This hour indicator is used to create a time stamp to enable CFR splitting and rejoining in order to circumvent the AAS maximum limitation of 30 systems entered on any one order. The quantity ordered for this feature is generated by eConfig.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 60 (Initial order maximum: 60)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9465) Qty Indicator

This quantity indicator is used to specify the remaining, or N-1 quantity of CFR entities that need to be accumulated for rejoining. The quantity ordered for this feature is generated by eConfig.and is equal to N-1, where 'N' equals the total quantity of CFRs being rejoined.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9466) Countable Member Indicator

This administrative indicator used to identify each CFR associated with a date/time stamp that is eligible for splitting and rejoining. The quantity ordered for this feature is generated by eConfig.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9700) Language Group Specify - Dutch

Dutch language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9703) Language Group Specify - French

French language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A

- Return parts MES: Does not apply

(#9704) Language Group Specify - German

German language group for Nomenclature and Standard Publications.

- Attributes provided: Language specify
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9705) Language Group Specify - Polish

Polish language group for Nomenclature and Standard Publications.

- Attributes provided: Language specify
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9706) Language Group Specify - Norwegian

Norwegian language group for Nomenclature and Standard Publications.

- Attributes provided: Language specify
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9707) Language Group Specify - Portuguese

Portuguese language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial

- CSU: N/A
- Return parts MES: Does not apply

(#9708) Language Group Specify - Spanish

Spanish language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9711) Language Group Specify - Italian

Italian language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9712) Language Group Specify - Canadian French

Canadian French language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9714) Language Group Specify - Japanese

Japanese language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None

- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9715) Language Group Specify - Traditional Chinese (Taiwan)

Traditional Chinese language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9716) Language Group Specify - Korean

Korean language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9718) Language Group Specify - Turkish

Turkish language group for nomenclature and publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9719) Language Group Specify - Hungarian

Hungarian language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)

- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9720) Language Group Specify - Slovakian

Slovakian language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9721) Language Group Specify - Russian

Russian language group for nomenclature and standard publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9722) Language Group Specify - Simplified Chinese (PRC)

Simplified Chinese language group for nomenclature and standard publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9724) Language Group Specify - Czech

Czech language group for nomenclature and standard publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0

- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9725) Language Group Specify - Romanian

Romanian language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9726) Language Group Specify - Croatian

Croatian language group for Nomenclature and Standard Publications.

- Attributes provided: Language specify
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9727) Language Group Specify - Slovenian

Slovenian language group for Nomenclature and Standard Publications.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9728) Language Group Specify - Brazilian Portuguese

Brazilian Portuguese language group for Nomenclature and Standard Publications.

- Attributes provided: Language specify
- Attributes required: None

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#9729) Language Group Specify - Thai

Thai language group for Nomenclature and Standard Publications.

- Attributes provided: Language specify
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#B0PR) SP Hard Drive/Media Retention - Power 5 years

This feature indicates ServicePac Hard Drive/Media Retention - Power 5 years

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum:)
- OS level required:
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#B0VH) SP HDR/MR POWER 3Y

ServicePac for Hard Drive or Media Retention for Power 3 years

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum:)
- OS level required:
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EB2J) 10m (30.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers

QSFP+ cable is used for 40Gb-to-40Gb Ethernet connectivity. Clients can use this QSFP+ Direct Attach Cable for Ethernet connectivity.

- Attributes provided: 10m QSFP+ to QSFP+ Cable
- Attributes required: QSFP/QSFP+ ports with optical transceivers

- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB2K) 30m (90.3-ft) - IBM MTP 12 strand cable for 40/100G transceivers

QSFP+ cable is used for 40Gb-to-40Gb Ethernet connectivity. Clients can use this QSFP+ Direct Attach Cable for Ethernet connectivity.

- Attributes provided: 30m QSFP+ to QSFP+ Cable
- Attributes required: QSFP/QSFP+ ports with optical transceivers
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB2Z) Lift Tool

A lift tool raises and lowers servers and I/O drawers so they can be placed into or removed from standard 19-inch racks. It allows heavier equipment to be handled more safely by fewer people. #EB27 has a hand crank to lift and position up to 350 pounds or 159 kg. The #EB2Z length and width are 44 inches x 24.5 inches or 1.12 meters x 0.62 meters. It has rollers which allow it to be moved to different racks in the data center.

- Attributes provided: Lift Tool
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB35) Mobile Enablement

Enterprise Pool Enablement for 9080-HEX

- Attributes provided: Enterprise Pool Enablement Code
- Attributes required: none
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#EB3Z) Lift tool based on GenieLift GL-8 (standard)

This feature delivers the Low-Cost Lift Tool (based on GenieLift GL-8 (standard)) for IBM servers.

Feature #EB3Z is a feature that is available on multiple Power8, Power9 and Power10 as well as the rack models 7965-S42, 7014-T00, and 7014-T42). Failure to have at least one Lift tool available in a location may result in delayed or prolonged maintenance times.

A lift tool raises and lowers servers and I/O drawers so they can be placed into or removed from standard 19-inch racks. It allows heavier equipment to be handled more safely by fewer people. Lift tool feature EB3Z has a hand crank to lift and position up to 181 kg (400 lbs). The lift tool feature EB3Z operating length and width are 88.3 cm x 62.9 cm (34 3/4 x 24 3/4 in). It has rollers which allow it to be moved to different racks in the data center.

- Attributes provided: Lift Tool
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EB3Z not orderable in Albania, Bahrain, Bulgaria, Croatia, Egypt, Greece, Jordan, Kuwait, Kosovo, Montenegro, Morocco, Oman, UAE, Qatar, Saudi Arabia, Serbia, Slovakia, Slovenia, Taiwan and Ukraine.

(#EB46) 10GbE Optical Transceiver SFP+ SR

One optical transceiver for 10Gb Ethernet adapter such as #EC2S or #EC2R and also #EC2U or #EC2T using SFP+ SP. Does not include cable.

The 10 Gb optical transceiver is capable up to 300 M through the OM3 cable or 82 M through OM2 cable. Either one or both of the adapter's two SFP+ ports can be populated.

- Attributes provided: Optical Transceiver SFP+ SR 10Gb
- Attributes required: SFP+ socket
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB47) 25GbE Optical Transceiver SFP28

One optical transceiver for 25Gb Ethernet adapter such as #EC2U or #EC2T using SFP28. Does not include cable.

The 25 Gb optical transceiver is capable up to 100 m through the OM4 cable or 70 M through OM3 cable. Either one or both of the adapters two SFP28 ports can be populated.
Note: The SFP28 25GbE transceiver only supports 25GbE speeds.

- Attributes provided: SFP28 optical transceiver
- Attributes required: SFP28 socket
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB48) 1GbE Base-T Transceiver RJ45

One 1GbE Base-T RJ45 transceiver for 25/10Gb Ethernet adapter such as features EC2U or EC2T. Does not include cable. The 1GbE Base-T RJ45 transceiver uses CAT5E STP cable up to 100m.

The transceiver can be plugged into the SFP28 port of EC2U or EC2T in order to provide a 1GbE Base-T RJ45 connection. Either one or both of the adapter's two SFP ports can be populated.

Limitations:

- SRIOV or RoCE capabilities are not supported on the EC2U or EC2T adapters using this feature EB48.
- Feature EB48 cannot be used with Feature EB46 - 10GbE Optical Transceiver SFP+ SR or EB47 - 25GbE Optical Transceiver SFP28.
- Attributes provided: 1Gb RJ45 transceiver
- Attributes required: SFP socket
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required:
  - IBM i supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB49) QSFP28 to SFP28 Connector

This feature provides a QSFP28 to SFP28 connector which enables connections between a single lane transceiver/cable and a quad-lane port. The QSA28 provides the option to connect an SFP28 transceiver or cable to a QSFP28 port 100 Gb/s switch or network card.

The adapter has a QSFP28 form factor with a receptacle for an SFP28 transceiver/AOC/DAC connector. The QSA28 interoperates with all major optical modules and direct attached copper cable. Its design assures minimum loss on the conversion path between the QSFP28 cage and the SFP28 receptacle.

The QSFP28 to SFP28 connector is qualified for 10GbE SFP+ and 1GbE SFP transceivers meeting the Small Form Factor Pluggable (SFP) Transceiver Multi-source Agreement (MSA).
Note: This feature is only available through Offering Management approval, contact douglasg@ca.ibm.com or bbarnett@us.ibm.com for authorization.

- Attributes provided: QSFP28 to SFP28 Adapter
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB4J) 0.5m SFP28/25GbE copper Cable

Feature EB4J is a passive 0,5 meter copper cable that contains a single high-speed copper pair, operating at data rates of up to 25 Gb/s. This cable can be used for either 25Gb Ethernet adapters or switches. Built onto each end of the cable is a passive SFP28 copper cable transceiver. Cables are available in various lengths: 0.5M - #EB4J, 1.0M=#EB4K, 1.5M=#EB4L, 2.0M=#EB4M

Attributes provided: Copper cable with SFP28 transceivers

-
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB4K) 1.0m SFP28/25GbE copper Cable

Feature EB4K is a passive 1.0 meter copper cable that contains a single high-speed copper pair, operating at data rates of up to 25 Gb/s. This cable can be used for either 25Gb Ethernet adapters or switches. Built onto each end of the cable is a passive SFP28 copper cable transceiver. Cables are available in various lengths: 0.5M - #EB4J, 1.0M=#EB4K, 1.5M=#EB4L, 2.0M=#EB4M

- Attributes provided: Copper cable with SFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB4M) 2.0m SFP28/25GbE copper Cable

Feature EB4M is a passive 2.0 meter copper cable that contains a single high-speed copper pair, operating at data rates of up to 25 Gb/s. This cable can be used for either 25Gb Ethernet adapters or switches. Built onto each end of the cable is a passive SFP28 copper cable transceiver. Cables are available in various lengths: 0.5M - #EB4J, 1.0M=#EB4K, 1.5M=#EB4L, 2.0M=#EB4M

- Attributes provided: Copper cable with SFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB4P) 2.0m QSFP28/100GbE copper split Cable to SFP28 4x25GbE

Feature EB4P is a 2.0 meter, active optical 100Gb E to 4x25Gb E splitter cable. It provides connectivity between system units with a QSFP28 port on one side and up to four different SFP28 ports on the other side, such as a switch and four servers.

- Attributes provided: Copper splitter cable with QSFP28 and 4x SFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB4Z) Service wedge shelf tool kit for EB3Z

This feature provides a separate tool kit to replace the flat shelf with a wedge/angle shelf at the client site.
Note: EB4Z wedge shelf is IBM SSR use only (due to safety labels/ instructions/certifications only for IBM and not filed for clients). A client can order feature EB4Z to ensure the tool is conveniently located on site in case an IBM SSR needed to use it and do not want to wait for the SSR to locate and bring in an EB4Z or to schedule additional personnel to manually handle server installation/removal from the rack.

Client is free to use EB3Z (without EB4Z) for their normal work.

- Attributes provided: Wedge/angle shelf
- Attributes required: Feature EB3Z
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No
  Note: Feature EB4Z not orderable in Albania, Bahrain, Bulgaria, Croatia, Egypt, Greece, Jordan, Kuwait, Kosovo, Montenegro, Morocco, Oman, UAE, Qatar, Saudi Arabia, Serbia, Slovakia, Slovenia, Taiwan and Ukraine.

(#EB57) QSFP+ 40GbE Base-SR4 Transceiver

IBM QSFP+ optical base-SR4 transceiver required for 40 Gbps ports which are not using copper QSFP+ transceiver.

The QSFP+ 40GbE Base-SR4 Transceiver is designed to deliver the following functions:

- Hot-swappable input/output device that plugs into a 40 Gigabit Ethernet QSFP port
- Interoperable with other IEEE-compliant 40GBASE interfaces where applicable
- High-speed electrical interface compliant to the IEEE 802.3ba standard
- QSFP Form factor, 2-wire I2C communication interface and other low-speed electrical interface compliant to SFF 8436 and QSFP Multisource Agreement (MSA)
- Attributes provided: QSFP+ transceiver for 40 Gbs ports
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB59) 100GbE Optical Transceiver QSFP28

One optical transceiver for 100Gb Ethernet adapter such as #EC3L or #EC3M or #EC66 or #EC67 using QSFP28. Does not include cable.

See also AOC fiber cables which include QSFP28 transceivers EB5R - EB5Y.

- Attributes provided: Optical Transceiver QSFP28 100Gb.
- Attributes required: Port on adapter with QSFP28 socket.

- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5K) 1.0M 100GbE Copper Cable QSFP28

1.0 meter length passive copper cable with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches. Copper cable is also called "copper twinax" or "DAC" (Direct Attach Copper).

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EB5M (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Copper cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5L) 1.5M 100GbE Copper Cable QSFP28

1.5 meter length passive copper cable with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches. Copper cable is also called "copper twinax" or "DAC" (Direct Attach Copper).

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EB5M (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Copper cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5M) 2.0M 100GbE Copper Cable QSFP28

2.0 meter length passive copper cable with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches. Copper cable is also called "copper twinax" or "DAC" (Direct Attach Copper).

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EB5M (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Copper cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
  OS level required: None

-
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5R) 3M 100GbE Optical Cable QSFP28 (AOC)

3 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5S) 5M 100GbE Optical Cable QSFP28 (AOC)

5 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5T) 10M 100GbE Optical Cable QSFP28 (AOC)

10 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
  CSU: Yes

-
- Return parts MES: No

(#EB5U) 15M 100GbE Optical Cable QSFP28 (AOC)

15 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5V) 20M 100GbE Optical Cable QSFP28 (AOC)

20 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5W) 30M 100GbE Optical Cable QSFP28 (AOC)

30 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5X) 50M 100GbE Optical Cable QSFP28 (AOC)

50 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB5Y) 100M 100GbE Optical Cable QSFP28 (AOC)

100 meter length Active Optical fiber Cable (AOC) with QSFP28 transceivers on each end. The cable used for 100Gb Ethernet adapters or switches.

Cables are available in various lengths. See shorter passive copper cables #EB5J - #EJM (0.5M - 2.0M) or see active optical fiber cables #EJ5R - #EJ5Y (3M - 100M).

- Attributes provided: Optical fiber cable with QSFP28 transceivers
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EB73) IBM i 7.3 Indicator

This feature is used to indicate the correct level of code when IBM i is specified.

- Attributes provided: IBM i 7.3 Indicator
- Attributes required: IBM i operating system
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required:
  - IBM i 7.3 supported
  - Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#EB74) IBM i 7.4 Indicator

This feature is used to indicate the correct level of code when IBM i is specified.

Attributes provided: IBM i 7.4 Indicator

-
- Attributes required: IBM i operating system
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required:
  - IBM i 7.4 supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EB75) - IBM i 7.5 Indicator

This feature is used to indicate the correct level of code when IBM i is specified.

- Attributes provided: IBM i 7.5 Indicator
- Attributes required: IBM i operating system
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required:
  - IBM i 7.5 supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EBAB) IBM Rack-mount Drawer Bezel and Hardware

Indicates that this order is for a rack-mount system requiring an IBM bezel and hardware.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EBAC) OEM Rack-mount Drawer Bezel and Hardware

Indicates that this order is for a rack-mount system requiring an OEM bezel and hardware.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EBAE) System Node (5U) Drawer Indicator for Solution Edition for Healthcare

This feature provides a 5U system node drawer. It includes the drawer, labels, 4 CPU Sockets, 32 DIMM Slots, 8 LOW PROFILE PCIE Slots.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 2 (Initial order maximum: 2)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EBK4) 1.6M USB Cable

1.6M USB CABLE, MALE-TO-MALE

- Attributes provided: 1.6M USB CABLE
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EC2R) PCIe3 LP 2-Port 10Gb NIC&ROCE SR/Cu Adapter

This PCIe Gen3 Ethernet adapter provides two 10 Gb SFP+ ports. The adapter supports both NIC and IBTA RoCE standards. RoCE is Remote Direct Memory Access (RDMA) over Converged Ethernet. Using RoCE, the adapter can support significantly greater bandwidth with low latency and minimize CPU overhead by more efficiently using memory access. This offloads the CPU from I/O networking tasks, improving performance and scalability.

Cables: For 10GbE, IBM offers Direct Attach (DAC) cables up to 5m. SFP-based transceivers are included on each end of the cables. See features EN01, EN02, EN03.

Transceivers: IBM qualifies and supports SFP+ optical transceiver (FC EB46) to install into the adapter. Customers can also use their own optical cabling and SFP+ optical transceiver for the other end. The 10 Gb optical transceiver is capable up to 300 M through the OM3 cable or 82 M through OM2 cable. Either one or both of the adapter's two SFP+ ports can be populated.

Feature code #EC2R and #EC2S have identical electronics and function and CCIN (58FA), but have different tail stock brackets. #EC2R is low profile and #EC2S is full high. The adapter is based on a Mellanox ConnectX-4 adapter which uses a ConnectX-4 Lx EN Network Controller.

Attributes:

- PCI Express 3.0 (up to 8GT/s) x8
- PCIe Gen 3.0 compliant, 1.1 and 2.0 compatible
- RDMA over Converged Ethernet (RoCE)
- NIC and RoCE are concurrently supported
- RoCE supported on Linux and AIX (7.2 and later)
- NIC supported on all OSes
- IEEE 802.3ae (10Gb Ethernet), IEEE 802.3ad (Link Aggregation & Failover), IEEE 802.3az (Energy Efficient Ethernet), IEEE 802.1Q/P (VLAN Tagging), IEEE 802.10au (Congestion Notification), IEEE 802.1Qbg, IEEE 802.3Qaz D0.2 (ETS), IEEE 802.1Qbb D1.0 (PFC), IEEE 1588v2 (PTP)

- Jumbo frame support up to 9.6KB
- VXLAN and NVGRE Overlay Network offload support
- TCP/UDP/IP stateless offload
- TCP checksum offload
- TCP segmentation offload
- UDP checksum offload
- MSI-X, MSI and support of legacy pin interrup
- NIM boot support
- PowerVM SR-IOV support

For SR-IOV FAQs visit https://community.ibm.com/community/user/power/ viewdocument/ sr-iov- faqs-test?CommunityKey=71e6bb8a-5b34-44da-be8b-277834a183b0&ta b=librarydocuments#Q3

- Attributes provided: 2-port 10Gb Ethernet
- Attributes required: Low profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - IBM i supported
  - AIX supported
  - SUSE Linux Enterprise Server supported with Mellanox OFED 5.4, or later
  - Red Hat Enterprise Linux supported with Mellanox OFED 5.4, or later
  - IBM i supports dedicated driver for the NIC function.
  - IBM i supports native SR-IOV for the NIC function.
  - IBM i 7.4 also adds support for dedicated ROCE and native SR-IOV ROCE, which is used only by IBM Db2 Mirror for i.
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported; NIC capability only.

(#EC2S) PCIe3 2-Port 10Gb NIC&ROCE SR/Cu Adapter

This PCIe Gen3 Ethernet adapter provides two 10 Gb SFP+ ports. The adapter supports both NIC and IBTA RoCE standards. RoCE is Remote Direct Memory Access (RDMA) over Converged Ethernet. Using RoCE, the adapter can support significantly greater bandwidth with low latency and minimize CPU overhead by more efficiently using memory access. This offloads the CPU from I/O networking tasks, improving performance and scalability.

Cables: For 10GbE, IBM offers Direct Attach (DAC) cables up to 5m. SFP-based transceivers are included on each end of the cables. See features EN01, EN02, EN03.

Transceivers: IBM qualifies and supports SFP+ optical transceiver (FC EB46) to install into the adapter. Customers can also use their own optical cabling and SFP+ optical transceiver for the other end. The 10 Gb optical transceiver is capable up to 300 M through the OM3 cable or 82 M through OM2 cable. Either one or both of the adapter's two SFP+ ports can be populated.

Feature code #EC2R and #EC2S have identical electronics and function and CCIN (58FA), but have different tail stock brackets. #EC2R is low profile and #EC2S is full high. The adapter is based on a Mellanox ConnectX-4 adapter which uses a ConnectX-4 Lx EN Network Controller.

Attributes:

- PCI Express 3.0 (up to 8GT/s) x8
- PCIe Gen 3.0 compliant, 1.1 and 2.0 compatible
- RDMA over Converged Ethernet (RoCE)
  NIC and RoCE are concurrently supported

-
- RoCE supported on Linux and AIX (7.2 and later)
- NIC supported on all OSes
- IEEE 802.3ae (10Gb Ethernet), IEEE 802.3ad (Link Aggregation & Failover), IEEE 802.3az (Energy Efficient Ethernet), IEEE 802.1Q/P (VLAN Tagging), IEEE 802.10au (Congestion Notification), IEEE 802.1Qbg, IEEE 802.3Qaz D0.2 (ETS), IEEE 802.1Qbb D1.0 (PFC), IEEE 1588v2 (PTP)
- Jumbo frame support up to 9.6KB
- VXLAN and NVGRE Overlay Network offload support
- TCP/UDP/IP stateless offload
- TCP checksum offload
- TCP segmentation offload
- UDP checksum offload
- MSI-X, MSI and support of legacy pin interrup
- NIM boot support
- PowerVM SR-IOV support

For SR-IOV FAQs visit https://community.ibm.com/community/user/power/ viewdocument/ sr-iov- faqs-test? CommunityKey=71e6bb8a-5b34-44da-be8b-277834a183b0&ta b=librarydocuments#Q3

- Attributes provided: 2-port 10Gb Ethernet Adapter
- Attributes required: Full high PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:
  - IBM i supported
  - AIX supported
  - Red Hat Enterprise Linux supported with Mellanox OFED 5.4, or later
  - SUSE Linux Enterprise Server supported with Mellanox OFED 5.4, or later
  - IBM i supports dedicated driver for the NIC function.
  - IBM i supports native SR-IOV for the NIC function.
  - IBM i 7.4 also adds support for dedicated ROCE and native SR-IOV ROCE, which is used only by IBM Db2 Mirror for i.
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported; NIC capability only.

(#EC2T) PCIe3 LP 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter

This PCIe Gen3 Ethernet adapter provides two 25/10 Gb SFP28 ports. The adapter supports both NIC and IBTA RoCE standards. RoCE is Remote Direct Memory Access (RDMA) over Converged Ethernet. Using RoCE, the adapter can support significantly greater bandwidth with low latency and minimize CPU overhead by more efficiently using memory access. This offloads the CPU from I/O networking tasks, improving performance and scalability.

Cables: For 25GbE, IBM offers SFP28 Passive Copper 25Gb Ethernet cables up to 2m. SFP28 based transceivers are included on each end of these cables. See features EB4J, EB4K, EB4L, and EB4M for a 0.5M, 1.0M, 1.5M and 2.0M copper cable.

For 10GbE, IBM offers Direct Attach (DAC) cables up to 5m. SFP-based transceivers are included on each end of the cables. See features EN01, EN02, EN03.

Transceivers: For 25 GbE, IBM qualifies and supports SFP28 optical transceiver (FC EB47) to install into the adapter. Customers can also use their own optical cabling and SFP28 optical transceiver for the other end. The 25 Gb optical transceiver is capable up to 100 m through the OM4 cable or 70 M through OM3 cable. Either one or both of the adapter's two SFP28 ports can be populated.

Note: The (FC EB47) SFP28 25GbE transceiver only supports 25GbE speeds.

For 10 GbE, IBM qualifies and supports SFP+ optical transceiver (FC EB46) to install into the adapter. Customers can also use their own optical cabling and SFP+ optical transceiver for the other end. The 10 Gb optical transceiver is capable up to 300 M through the OM3 cable or 82 m through OM2 cable. Either one or both of the adapter's two SFP28 ports can be populated.

For 1GbE, IBM qualifies and supports the 1GbE Base-T Transceiver RJ45 (FC EB48) for 25/10Gb Ethernet adapter (features EC2U or EC2T). The 1GbE Base-T RJ45 transceiver uses CAT5E STP cable up to 100m (does not include cable). The transceiver can be plugged into the SFP28 port of the adapter in order to provide a 1GbE Base-T RJ45 connection.

Feature code #EC2T and #EC2U have identical electronics and function and CCIN (58FB), but have different tail stock brackets. #EC2T is low profile and #EC2U is full high. The adapter is based on a Mellanox ConnectX-4 adapter which uses a ConnectX-4 Lx EN Network Controller.

Attributes:

- PCI Express 3.0 (up to 8GT/s) x8
- PCIe Gen 3.0 compliant, 1.1 and 2.0 compatible
- RDMA over Converged Ethernet (RoCE)
- NIC and RoCE are concurrently supported
- RoCE supported on Linux and AIX (7.2 and later)
- NIC supported on all OSes
- IEEE 802.3ae (25Gb or 10Gb Ethernet), IEEE 802.3ad (Link Aggregation & Failover), IEEE 802.3az (Energy Efficient Ethernet), IEEE 802.1Q/P (VLAN Tagging), IEEE 802.10au (Congestion Notification), IEEE 802.1Qbg, IEEE 802.3Qaz D0.2 (ETS), IEEE 802.1Qbb D1.0 (PFC), IEEE 1588v2 (PTP)
- Jumbo frame support up to 9.6KB
- VXLAN and NVGRE Overlay Network offload support
- TCP/UDP/IP stateless offload
- TCP checksum offload
- TCP segmentation offload
- UDP checksum offload
- MSI-X, MSI and support of legacy pin interrup
- NIM boot support
- PowerVM SR-IOV support

For SR-IOV FAQs visit https://community.ibm.com/community/user/power/ viewdocument/ sr-iov- faqs-test? CommunityKey=71e6bb8a-5b34-44da-be8b-277834a183b0&ta b=librarydocuments#Q3

- Attributes provided: 2-port 25/10Gb Ethernet
- Attributes required: Low profile PCIe Gen3 slot

(#EC2U) PCIe3 2-Port 25/10Gb NIC&ROCE SR/Cu Adapter

This PCIe Gen3 Ethernet adapter provides two 25/10 Gb SFP28 ports. The adapter supports both NIC and IBTA RoCE standards. RoCE is Remote Direct Memory Access (RDMA) over Converged Ethernet. Using RoCE, the adapter can support significantly greater bandwidth with low latency and minimize CPU overhead by more efficiently using memory access. This offloads the CPU from I/O networking tasks, improving performance and scalability.

Cables: For 25GbE, IBM offers SFP28 Passive Copper 25Gb Ethernet cables up to 2m. SFP28 based transceivers are included on each end of these cables. See features EB4J, EB4K, EB4L, and EB4M for a 0.5M, 1.0M, 1.5M and 2.0M copper cable.

For 10GbE, IBM offers Direct Attach (DAC) cables up to 5m. SFP-based transceivers are included on each end of the cables. See features EN01, EN02, EN03.

Transceivers: For 25 GbE, IBM qualifies and supports SFP28 optical transceiver (FC EB47) to install into the adapter. Customers can also use their own optical cabling and SFP28 optical transceiver for the other end. The 25 Gb optical

transceiver is capable up to 100 m through the OM4 cable or 70 M through OM3 cable. Either one or both of the adapter's two SFP28 ports can be populated.
Note: The (FC EB47) SFP28 25GbE transceiver only supports 25GbE speeds.

For 10 GbE, IBM qualifies and supports SFP+ optical transceiver (FC EB46) to install into the adapter. Customers can also use their own optical cabling and SFP+ optical transceiver for the other end. The 10 Gb optical transceiver is capable up to 300 M through the OM3 cable or 82 m through OM2 cable. Either one or both of the adapter's two SFP28 ports can be populated.

For 1GbE, IBM qualifies and supports the 1GbE Base-T Transceiver RJ45 (FC EB48) for 25/10Gb Ethernet adapter (features EC2U or EC2T). The 1GbE Base-T RJ45 transceiver uses CAT5E STP cable up to 100m (does not include cable). The transceiver can be plugged into the SFP28 port of the adapter in order to provide a 1GbE Base-T RJ45 connection.

Feature code #EC2T and #EC2U have identical electronics and function and CCIN (58FB), but have different tail stock brackets. #EC2T is low profile and #EC2U is full high. The adapter is based on a Mellanox ConnectX-4 adapter which uses a ConnectX-4 Lx EN Network Controller.

Attributes:

- PCI Express 3.0 (up to 8GT/s) x8
- PCIe Gen 3.0 compliant, 1.1 and 2.0 compatible
- RDMA over Converged Ethernet (RoCE)
- NIC and RoCE are concurrently supported
- RoCE supported on Linux and AIX (7.2 and later)
- NIC supported on all OSes
- IEEE 802.3ae (25Gb or 10Gb Ethernet), IEEE 802.3ad (Link Aggregation & Failover), IEEE 802.3az (Energy Efficient Ethernet), IEEE 802.1Q/P (VLAN Tagging), IEEE 802.10au (Congestion Notification), IEEE 802.1Qbg, IEEE 802.3Qaz D0.2 (ETS), IEEE 802.1Qbb D1.0 (PFC), IEEE 1588v2 (PTP)
- Jumbo frame support up to 9.6KB
- VXLAN and NVGRE Overlay Network offload support
- TCP/UDP/IP stateless offload
- TCP checksum offload
- TCP segmentation offload
- UDP checksum offload
- MSI-X, MSI and support of legacy pin interrup
- NIM boot support
- PowerVM SR-IOV support
  For SR-IOV FAQs visit https://community.ibm.com/community/user/power/viewdocument/sr-iov-faq s-t est?CommunityKey=71e6bb8a-5b34-44da-be8b-277834a183b0&tab=librarydocum ent s#Q3

- Attributes provided: 2-port 25/10Gb Ethernet Adapter
- Attributes required: Full high PCIe Gen3 slot

(#EC5J) Mainstream 800 GB SSD NVMe U.2 module

Mainstream 800 GB solid state drive (SSD) formatted in 4096 byte sectors (4K). The drive is U.2 2.5" 7mm form factor supporting NVMe PCIe Gen3 x 4 interface. Recommend to be used for boot support and non-write intensive workloads.

DWPD (Drive Write Per Day) rating is 2.4 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (PB)
800 GB	3.5
1.6 TB	7
3.2 TB	14

Depending on the nature of the workload, the lifetime TBW may be somewhat larger.

To read the warranty and maintenance applicable to mainstream devices on POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

- Attributes provided: 800 GB low latency flash memory
- Attributes required: PCIe Gen3 slot in system unit
- Minimum required: 0
- Maximum allowed: 16 (Initial order maximum: 16)
- OS level required:
  - AIX - supported
  - IBM i - not supported
  - IBM i - not supported via VIOS
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EC5K) Mainstream 1.6 TB SSD NVMe U.2 module

Mainstream 1.6 TB solid state drive (SSD) formatted in 4096 byte sectors (4K). The drive is U.2 2.5" 7mm form factor supporting NVMe PCIe Gen3 x 4 interface. Recommend to be used for boot support and non-write intensive workloads.

DWPD (Drive Write Per Day) rating is 2.4 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (PB)
800 GB	3.5
1.6 TB	7
3.2 TB	14

Depending on the nature of the workload, the lifetime TBW may be somewhat larger.

To read the warranty and maintenance applicable to mainstream devices on POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

- Attributes provided: 1.6 TB low latency flash memory
- Attributes required: PCIe Gen3 slot in system unit
- Minimum required: 0
- Maximum allowed: 16 (Initial order maximum: 0)
- OS level required:
  - IBM i - not supported and IBM i not supported via VIOS
  - AIX - supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EC5L) Mainstream 3.2 TB SSD NVMe U.2 module

Mainstream 3.2 TB solid state drive (SSD) formatted in 4096 byte sectors (4K). The drive is U.2 2.5" 7mm form factor supporting NVMe PCIe Gen3 x 4 interface. Recommend to be used for boot support and non-write intensive workloads.

DWPD (Drive Write Per Day) rating is 2.4 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (PB)

800 GB 3.5
1.6 TB 7
3.2 TB 14

Depending on the nature of the workload, the lifetime TBW may be somewhat larger.

To read the warranty and maintenance applicable to mainstream devices on POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

- Attributes provided: 3.2 TB low latency flash memory
- Attributes required: PCIe Gen3 slot in system unit
- Minimum required: 0
- Maximum allowed: 16 (Initial order maximum: 0)
- OS level required:
  - IBM i - not supported and IBM i not supported via VIOS
  - AIX - supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EC67) PCIe4 LP 2-port 100Gb ROCE EN LP adapter

This PCIe Gen4 Ethernet x16 adapter provides two 100 GbE QSFP28 ports. The adapter supports both NIC and IBTA RoCE standards. RoCE is Remote Direct Memory Access (RDMA) over Converged Ethernet. Using RoCE, the adapter can support significantly greater bandwidth with low latency and minimize CPU overhead by more efficiently using memory access. This offloads the CPU from I/O networking tasks, improving performance and scalability.

For 100GbE network connectivity, IBM offers either passive copper twinax cables up to 2 meters in length or active optical cables up to 100 meters in length. See features #EB5J - #EB5M for a 0.5M, 1.0M, 1.5M and 2.0M copper cable. See features #EB5R - #EB5Y for a 3M, 5M, 10M, 15M, 20M, 30M, 50M or 100M active optical cable. Transceivers are included on each end of these QSFP28 cables. Alternatively to the above supported cables, you may choose to order an IBM qualified and supported QSFP28 optical transceiver (feature #EB59) to put into the adapter ports and MTP/MPO cable 10M or 30M in length (feature #EB2J or #EB2K) optical cabling to use with your own QSP28 optical transceiver for the other end.

For 40GbE network connectivity, IBM offers either passive copper twinax cables up to 3 meters in length. See features #EB2B and EB2H for a 1.0M and 3.0M copper cable. Transceivers are included on each end of these QSFP+ cables. Alternatively to the above supported cables, you may choose to order an IBM qualified and supported QSFP+ 40G BASE-SR optical transceiver (feature #EB27) to put into the adapter ports and MTP/MPO cable 10M or 30M in length (feature #EB2J or #EB2K) optical cabling to use with your own QSP28 optical transceiver for the other end.

Either one or both of the adapter's two QSP28 ports can be populated. When two ports are filled, both can have copper cables, both can have optical cables, or one can be copper and one can be optical.

Feature code #EC66 and #EC67 have identical electronics and function and CCIN (2CF3), but have different tail stock brackets. #EC66 is high profile and #EC67 is low profile. The adapter is based on a Mellanox ConnectX-5 adapter which uses a ConnectX-5 EN Network Controller.

Attributes:

- PCI Express 4.0 (up to 16GT/s) x16
- PCIe Gen 4.0 compliant, 1.1, 2.0 and 3.0 compatible
- RDMA over Converged Ethernet (RoCE)
- NIC and RoCE are concurrently supported
- RoCE supported on Linux and AIX
- NIC supported on all OSes
- TCP/UDP/IP stateless offload

- LSO, LRO, checksum offload
- NIM boot support
- Backward compatible with 40Gb Ethernet when using compatible cables/transceivers
- PowerVM SR-IOV support

For SR-IOV FAQs visit https://community.ibm.com/community/user/power/ viewdocument/ sr-iov- faqs-test? CommunityKey=71e6bb8a-5b34-44da-be8b-277834a183b0&ta b=librarydocuments#Q3

- Attributes provided: Dual-port 100 GbE Adapter with RoCE capability
- Attributes required: Available Gen4 PCIe Slot
- Minimum required: 0
- Maximum allowed: 24 (Initial order maximum: 24)
- OS level required:
  - AIX - supported
  - IBM i supported
  - SUSE Linux Enterprise Server supported with Mellanox OFED 5.4, or later
  - Red Hat Enterprise Linux supported with Mellanox OFED 5.4, or later
  - IBM i supports dedicated driver for the NIC function.
  - IBM i supports native SR-IOV for the NIC function.
  - IBM i 7.4 also adds support for dedicated ROCE and native SR-IOV ROCE, which is used only by IBM Db2 Mirror for i. 'IBM i native SR-IOV'.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported; NIC capability only.

(#EC6J) PCIe2 LP 2-Port USB 3.0 Adapter

The PCIe Gen2 x8 short 2-port USB 3.0 adapter provides support for USB devices. The #EC6J and #EC6K USB adapters are electronically identical with the same 590F CCIN. They differ physically in their tailstock. #EC6J is low profile and #EC6K is full high.

- Attributes provided: Connectivity with USB 2.0 - 3.0 capable devices
- Attributes required: One low profile available PCIe slot
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - AIX - supported
  - Linux - supported
  - IBM i - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EC6N) E1080 Control Unit External USB Enablement Cable

1.6M USB CABLE, MALE-TO-MALE

- Attributes provided: 1.6M Internal USB Cable
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)

- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EC77) PCIe4 LP 2-port 100Gb Crypto Connectx-6 DX QFSP56

This Gen4 x16 adapter provides up to two ports of 100 Gb/s Ethernet connectivity. The adapter supports both the network interface card (NIC) standard and the InfiniBand trade association (IBTA) standard for remote direct memory access (RDMA) over converged Ethernet (RoCE). Using RoCE, the adapter can support significantly greater bandwidth with low latency and minimize CPU overhead by more efficiently using memory access. This offloads the CPU from I/O networking tasks, improving performance and scalability.

Feature EC77 and feature EC78 are sourced from the Mellanox Corporation and are based on the Mellanox ConnectX-6 Dx technology. They have identical electronics and function and the same CCIN of 2CFA, but they have different tailstocks. Feature EC77 is low profile and feature EC78 is full height.

- Limitations:
  - The adapter is supported only by IBM i and runs in dedicated mode only (no PowerVM virtualization) and RoCE and IP Security (IPSEC). (RoCE and IPSEC are only supported by IBM i Db2 Mirror).
  - Single root I/O virtualization (SR-IOV) is not supported.
- Details for the ports include:
  - PCI Express Gen4 x16
  - IPSEC over RoCE
  - RoCE over overlay networks
  - NIC supported
  - Attributes provided: Two-port 100Gb/s Ethernet adapter
  - Attributes required: One low-profile PCIe slot Gen4 x16
  - Minimum required: 0
  - Maximum allowed: 24 (Initial order maximum: 24)
  - OS level required:
    - IBM i supported

    Note: IBM i 7.3 supports dedicated driver for the NIC function.
    Note: IBM i 7.4 supports dedicated driver for the NIC, RoCE, and IPSEC functions (RoCE and IPSEC are only supported by IBM i Db2 Mirror).
  - Initial Order/MES/Both/Supported: Both
  - CSU: Yes
  - Return parts MES: No

(#EC7Q) 800GB Mainstream NVMe U.2 SSD 4k for AIX/Linux

Mainstream 800 GB NVMe SFF U.2 7mm SSD (Solid State Drive) is a Peripheral Component Interconnect Express (PCIe) generation 4 (Gen4) drive. The SSD can be used in any U.2 7mm Non-Volatile Memory Express (NVMe) PCIe Gen4 x 4 interface in the system. NVMe is a high performance architecture and command protocol that can read/write flash memory. Compared to a SAS or SATA SSD, the NVMe SSD provides more read/write input/output operations per second (IOPS) and larger throughput (GB/sec).

Feature #EC7Q has the CCIN number of 59B4 and indicates usage by AIX, Linux, or VIOS in which the SSD is formatted in 4096 byte sectors.

Mainstream 800 GB solid state drive (SSD) formatted in 4096 byte sectors (4K). This Mainstream solid-state drive is rated at 2.4 DWPD (Drive Writes Per Day) calculated over a 5-year period for 100% (4K bytes or larger) random write workloads. Use for workloads within this rating are fully supported and will maintain high reliability and MTBF.

The nature of the workload has a great impact on the maximum write capacity. If a high percentage of more sequentially oriented writes is used instead of random writes, the maximum write capacity will be larger. Writes past the SSD's maximum write capacity will continue to work for some period of time, but perform much more slowly. Whether the

application uses sequential or random reads from the device does not affect the life of the device. A Predictive Failure Analysis message will indicate that it is time to replace the SSD if enabled by the operating system (OS) and system administrator. Customers are recommended to monitor SMART log critical information via the appropriate OS utility to observe drive life remaining information. IBM NVMe SSD failures will be replaced during the standard warranty and maintenance period for SSD's that have not reached the maximum number of write cycles. SSD's that reach this limit may fail to operate according to specifications and must be replaced at the client's expense. Boot supported. Data redundancy on a failed SSD may be provided by OS mirroring or software RAID wherever applicable.

- Attributes provided: 800 GB low latency flash memory
- Attributes required: PCIe Gen3 slot in system unit
- Minimum required: 0
- Maximum allowed: 16 (Initial order maximum: 16)
- OS level required:
  - AIX supported
  - IBM i not supported
  - IBM i not supported via VIOS
  - 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 8.4, for Power LE, or later
  - Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8.4, for Power LE, or later
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBJ) SAS X Cable 3m - HD Narrow 6Gb 2-Adapters to Enclosure

This 3 meter SAS cable connects two PCIe2 SAS adapters or two PCIe3 SAS adapters to a SAS I/O enclosure. This X cable has four connectors, two Mini-SAS HD (High Density) Narrow connector and two Mini SAS connectors. The Mini-SAS HD Narrow connectors attach to two SAS adapters such as two #EJ0J or two #EJ0L or two #5913. Both Mini SAS connectors attach to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

This cable is almost identical to the #3454 3m SAS X cable, except #ECBJ connectors for the SAS adapters are more narrow allowing it to fit onto the newer PCIe3 SAS adapters.

- Attributes provided: Connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS I/O drawer with Mini-SAS connectors or between PCIe2 SAS adapter with Mini-SAS HD connectors and a SAS I/O drawer with Mini-SAS connectors.
- Attributes required: available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBK) SAS X Cable 6m - HD Narrow 6Gb 2-Adapters to Enclosure

This 3 meter SAS cable connects two PCIe2 SAS adapters or two PCIe3 SAS adapters to a SAS I/O enclosure. This X cable has four connectors, two Mini-SAS HD (High Density) Narrow connector and two Mini SAS connectors. The Mini-SAS HD Narrow connectors attach to two SAS adapters such as two #EJ0J or two #EJ0L or two #5913. Both Mini SAS connectors attach to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This

cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

This cable is almost identical to the #3454 3m SAS X cable, except #ECBJ connectors for the SAS adapters are more narrow allowing it to fit onto the newer PCIe3 SAS adapters.

This 6 meter SAS cable connects two SAS adapters to a SAS I/O enclosure. This X cable has four connectors, two Mini-SAS HD (High Density) connector and two Mini SAS connectors. The Mini-SAS HD connectors attaches to two SAS adapters such as the #5913 1.8GB RAID SAS Adapter. Both Mini SAS connectors attach to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapters can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: When this cable is ordered with a system in a rack specifying IBM Plant integration, IBM Manufacturing will ship SAS cables longer than 3 meters in a separate box and not attempt to place the cable in the rack.

- Attributes provided: Connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS I/O drawer with Mini-SAS connectors or between PCIe2 SAS adapter with Mini-SAS HD connectors and a SAS I/O drawer with Mini-SAS connectors.
- Attributes required: available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBM) SAS X Cable 15m - HD Narrow 3Gb 2-Adapters to Enclosure

This 15 meter SAS cable connects two PCIe2 SAS adapters or two PCIe3 SAS adapters to a SAS I/O enclosure. This X cable has four connectors, two Mini-SAS HD (High Density) Narrow connector and two Mini SAS connectors. The Mini-SAS HD Narrow connectors attach to two SAS adapters such as two #EJ0J or two #EJ0L or two #5913. Both Mini SAS connectors attach to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

This cable is almost identical to the #3458 15m SAS X cable, except #ECBM connectors for the SAS adapters are more narrow allowing it to fit onto the newer PCIe3 SAS adapters.

Limitation: When this cable is ordered with a system in a rack specifying IBM Plant integration, IBM Manufacturing will ship SAS cables longer than 3 meters in a separate box and not attempt to place the cable in the rack.

- Attributes provided: Connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS I/O drawer with Mini-SAS connectors or between PCIe2 SAS adapter with Mini-SAS HD connectors and a SAS I/O drawer with Mini-SAS connectors.
- Attributes required: available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported

- CSU: Yes
- Return parts MES: No

(#ECBT) SAS YO Cable 1.5m - HD Narrow 6Gb Adapter to Enclosure

This 1.5 meter SAS cable connects a PCIe2 SAS adapter or a PCIe3 SAS adapter to a SAS I/O enclosure. This YO cable has three connectors, one Mini-SAS HD (High Density) Narrow connector and two Mini SAS connectors. The Mini-SAS HD Narrow connector attaches to a SAS adapter such as the #EJ0J or #EJ0L or #5913. Both Mini SAS connectors attach to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable.

This cable is almost identical to the #3450 1.5m SAS YO cable, except the #ECBT connector for the SAS adapters is more narrow allowing it to fit onto the newer PCIe3 SAS adapters.

- Attributes provided: Connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS I/O drawer with Mini-SAS connectors or between PCIe2 SAS adapter with Mini-SAS HD connectors and a SAS I/O drawer with Mini-SAS connectors.
- Attributes required: available connectors on SAS controller and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBU) SAS YO Cable 3m - HD Narrow 6Gb Adapter to Enclosure

This 3 meter SAS cable connects a PCIe2 SAS adapter or a PCIe3 SAS adapter to a SAS I/O enclosure. This YO cable has three connectors, one Mini-SAS HD (High Density) Narrow connector and two Mini SAS connectors. The Mini-SAS HD Narrow connector attaches to a SAS adapter such as the #EJ0J or #EJ0L or #5913. Both Mini SAS connectors attach to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable.

This cable is almost identical to the #3451 3m SAS YO cable, except the #ECBU connector for the SAS adapters is more narrow allowing it to fit onto the newer PCIe3 SAS adapters.

- Attributes provided: Connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS I/O drawer with Mini-SAS connectors or between PCIe2 SAS adapter with Mini-SAS HD connectors and a SAS I/O drawer with Mini-SAS connectors.
- Attributes required: available connectors on SAS controller and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBV) SAS YO Cable 6m - HD Narrow 6Gb Adapter to Enclosure

This 6 meter SAS cable connects a PCIe2 SAS adapter or a PCIe3 SAS adapter to a SAS I/O enclosure. This YO cable has three connectors, one Mini-SAS HD (High Density) Narrow connector and two Mini SAS connectors. The Mini-SAS HD Narrow connector attaches to a SAS adapter such as the #EJ0J or #EJ0L or #5913. Both Mini SAS connectors attach

to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable.

This cable is almost identical to the #3452 6m SAS YO cable, except the #ECBV connector for the SAS adapters is more narrow allowing it to fit onto the newer PCIe3 SAS adapters.

Limitation: When this cable is ordered with a system in a rack specifying IBM Plant integration, IBM Manufacturing will ship SAS cables longer than 3 meters in a separate box and not attempt to place the cable in the rack.

- Attributes provided: Connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS I/O drawer with Mini-SAS connectors or between PCIe2 SAS adapter with Mini-SAS HD connectors and a SAS I/O drawer with Mini-SAS connectors.
- Attributes required: available connectors on SAS controller and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBW) SAS YO Cable 10m - HD Narrow 6Gb Adapter to Enclosure

This 10 meter SAS cable connects a PCIe2 SAS adapter or a PCIe3 SAS adapter to a SAS I/O enclosure. This YO cable has three connectors, one Mini-SAS HD (High Density) Narrow connector and two Mini SAS connectors. The Mini-SAS HD Narrow connector attaches to a SAS adapter such as the #EJ0J or #EJ0L or #5913. Both Mini SAS connectors attach to the same I/O drawer (enclosure) such as a #5886 EXP12S or #5887 EXP24S SAS I/O drawer. This cable can support up to 6Gb throughput.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/O drawer. See SAS cabling documentation for the length of each leg of the cable.

This cable is almost identical to the #3453 10m SAS YO cable, except the #ECBW connector for the SAS adapters is more narrow allowing it to fit onto the newer PCIe3 SAS adapters.

Limitation: When this cable is ordered with a system in a rack specifying IBM Plant integration, IBM Manufacturing will ship SAS cables longer than 3 meters in a separate box and not attempt to place the cable in the rack.

- Attributes provided: Connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS I/O drawer with Mini-SAS connectors or between PCIe2 SAS adapter with Mini-SAS HD connectors and a SAS I/O drawer with Mini-SAS connectors.
- Attributes required: available connectors on SAS controller and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBY) SAS AE1 Cable 4m - HD Narrow 6Gb Adapter to Enclosure

This 4 meter SAS cable connects a PCIe3 SAS adapter to a SAS tape drive or DVD. The tape drive or DVD is probably in an I/O enclosure such as a bridge box or 1U media enclosure or tape library. This AE cable has two connectors, one Mini-SAS HD (High Density) Narrow connector and one Mini-SAS connectors. The Mini-SAS HD Narrow connector attaches to a SAS adapter such as the #EJ0X. The Mini-SAS connector attaches to a SAS tape drive enclosure or DVD. This cable can support up to 6Gb throughput.

Use #ECBY when ordering the cable as a feature code on a Power System. Alternatively the same cable can be ordered using feature code #5507 of the IBM tape enclosure or DVD.

- Attributes provided: connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and a SAS tape drive or DVD with Mini-SAS connectors
- Attributes required: available connectors on SAS controller such as #EJ0X, #EJ10 or #EJ11 for use with an available SAS tape drive or DVD.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECBZ) SAS YE1 Cable 3m - HD Narrow 6Gb Adapter to Enclosure

This 3 meter SAS cable connects a PCIe3 SAS adapter to one or two SAS tape drives. The tape drive(s) is probably in an I/O enclosure such as a bridge box or 1U media enclosure or tape library. This YO cable has three connectors, one Mini-SAS HD (High Density) Narrow connector and two Mini-SAS connectors. The Mini-SAS HD Narrow connector attaches to a SAS adapter such as the #EJ0X. Each Mini-SAS connector attaches to a different SAS tape drive enclosure. This cable can support up to 6Gb throughput.

Use #ECBZ when ordering the cable as a feature code on a Power System. Alternatively the same cable can be ordered using feature code #5509 of the IBM tape enclosure.

- Attributes provided: connection between PCIe3 SAS adapter with Mini-SAS HD Narrow connectors and one or two SAS tape drives with Mini-SAS connectors.
- Attributes required: available connectors on SAS controller such as #EJ0X, #EJ10 or #EJ11 for use with an available SAS tape drive.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECCR) 2M Active Optical Cable Pair for PCIe3 Expansion Drawer

The 2.0 meter active optical cable (AOC) pair connects a PCIe3 module in the PCIe Gen3 I/O Expansion Drawer to a PCIe3 Optical Converter Adapter in the system unit. There are two identical cables in the cable pair, each with two CXP connectors. One of the cables attaches to the top CXP port of the PCIe3 module and to the top CXP port of the PCIe3 Optical Converter Adapter. The other cable attaches to the bottom CXP ports. See also other AOC cable length options such as the feature feature #ECCY (10 meter) or #ECCZ (20 meter). The 2 meter length is suggested only for intra-rack connections.

- Attributes provided: Pair of 2 meter active optical cables
- Attributes required: CXP ports on a PCIe3 Optical Cable Adapter and on a PCIe3 module in a PCIe Gen3 Expansion Drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECCY) 10M Active Optical Cable Pair for PCIe3 Expansion Drawer

The 10.0 meter active optical cable (AOC) pair connects a PCIe3 module in the PCIe Gen3 I/O Expansion Drawer to a PCIe3 Optical Converter Adapter in the system unit. There are two identical cables in the cable pair, each with two CXP connectors. One of the cables attaches to the top CXP port of the PCIe3 module and to the top CXP port of the PCIe3 Optical Converter Adapter. The other cable attaches to the bottom CXP ports.

- Attributes provided: Pair of 10 meter active optical cables
- Attributes required: CXP ports on a PCIe3 Optical Cable Adapter and on a PCIe3 module in a PCIe Gen3 Expansion Drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECCZ) 20M Active Optical Cable Pair for PCIe3 Expansion Drawer

The 20 meter active optical cable (AOC) pair connects a PCIe3 module in the PCIe Gen3 I/O Expansion Drawer to a PCIe3 Optical Converter Adapter in the system unit. There are two identical cables in the cable pair, each with two CXP connectors. One of the cables attaches to the top CXP port of the PCIe3 module and to the top CXP port of the PCIe3 Optical Converter Adapter. The other cable attaches to the bottom CXP ports. See also other AOC cable length options such as the feature #ECCR (2 meter) or feature #ECCY (10 meter). The 20 meter length is suggested only for extra-rack connections.

- Attributes provided: Pair of 20 meter active optical cables
- Attributes required: CXP ports on a PCIe3 Optical Cable Adapter and on a PCIe3 module in a PCIe Gen3 Expansion Drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECDJ) 3.0M SAS X12 Cable (Two Adapter to Enclosure)

This 3 meter SAS cable connects two SAS adapters to a SAS I/O enclosure. This X cable has four Mini-SAS HD (High Density) connectors. Two of connectors attach to two SAS adapters such as in the PCIe3 SAS Adapter. The other two connectors attach to one SAS I/ O Enclosure such as the EXP12SX (#ESLL/#ELLL) or EXP24SX (#ESLS/ #ELLS). This cable is designed for high speed (DHS) to support up to 12Gb throughput if the adapter has that capability.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/ O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: This cable cannot be used with the EXP24S I/O drawer (#5887 or# EL1S) which uses Mini-SAS connectors which are not HD.

- Attributes provided: Connection between two SAS adapters with Mini-SAS HD connectors and a SAS I/O drawer with Mini SAS HD connectors
- Attributes required: Available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None

  Initial Order/MES/Both/Supported: Both

-
- CSU: Yes
- Return parts MES: No

(#ECDK) 4.5M SAS X12 Active Optical Cable (Two Adapter to Enclosure)

This 4.5 meter SAS cable connects two SAS adapters to a SAS I/O enclosure. This X cable has four Mini-SAS HD (High Density) connectors. Two of connectors attach to two SAS adapters such as in the PCIe3 SAS Adapter. The other two connectors attach to one SAS I/ O Enclosure such as the EXP12SX (#ESLL/#ELLL) or EXP24SX (#ESLS/ #ELLS). This cable is designed for high speed (DHS) to support up to 12Gb throughput if the adapter has that capability.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/ O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: This cable cannot be used with the EXP24S I/O drawer (#5887 or# EL1S) which uses Mini-SAS connectors which are not HD.
Note: AOC cables require minimum level of SAS Adapter firmware. Refer to the latest prerequisites at

https://www14.software.ibm.com/webapp/set2/iprt/home

- Attributes provided: Connection between two SAS adapters with Mini-SAS HD connectors and a SAS I/O drawer with Mini SAS HD connectors
- Attributes required: Available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECDL) 10M SAS X12 Active Optical Cable (Two Adapter to Enclosure)

This 10 meter SAS cable connects two SAS adapters to a SAS I/O enclosure. This X cable has four Mini-SAS HD (High Density) connectors. Two of connectors attach to two SAS adapters such as in the PCIe3 SAS Adapter. The other two connectors attach to one SAS I/ O Enclosure such as the EXP12SX (#ESLL/#ELLL) or EXP24SX (#ESLS/ #ELLS). This cable is designed for high speed (DHS) to support up to 12Gb throughput if the adapter has that capability.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/ O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: This cable cannot be used with the EXP24S I/O drawer (#5887 or# EL1S) which uses Mini-SAS connectors which are not HD.
Note: AOC cables require minimum level of SAS Adapter firmware. Refer to the latest prerequisites at

https://www14.software.ibm.com/webapp/set2/iprt/home

- Attributes provided: Connection between two SAS adapters with Mini-SAS HD connectors and a SAS I/O drawer with Mini SAS HD connectors
- Attributes required: Available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECDT) 1.5M SAS YO12 Cable (Adapter to Enclosure)

This 1.5 meter SAS cable connects one SAS adapter to a SAS I/O enclosure. This YO cable has three Mini-SAS HD (High Density) connectors. One of connectors attach to a SAS adapter such as in the PCIe3 SAS Adapter. The other two connectors attach to one SAS I/O Enclosure such as the EXP12SX (#ESLL/#ELLL) or EXP24SX (#ESLS/ #ELLS). This cable is designed for high speed (DHS) to support up to 12Gb throughput if the adapter has that capability.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/ O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: This cable cannot be used with the EXP24S I/O drawer (#5887 or# EL1S) which uses Mini-SAS connectors which are not HD.

- Attributes provided: Connection between a SAS adapters with Mini-SAS HD connectors and a SAS I/O drawer with Mini SAS HD connectors
- Attributes required: Available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECDU) 3.0M SAS YO12 Cable (Adapter to Enclosure)

This 3 meter SAS cable connects one SAS adapter to a SAS I/O enclosure. This YO cable has three Mini-SAS HD (High Density) connectors. One of connectors attach to a SAS adapter such as in the PCIe3 SAS Adapter. The other two connectors attach to one SAS I/O Enclosure such as the EXP12SX (#ESLL/#ELLL) or EXP24SX (#ESLS/ #ELLS). This cable is designed for high speed (DHS) to support up to 12Gb throughput if the adapter has that capability.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/ O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: This cable cannot be used with the EXP24S I/O drawer (#5887 or# EL1S) which uses Mini-SAS connectors which are not HD.

- Attributes provided: Connection between a SAS adapters with Mini-SAS HD connectors and a SAS I/O drawer with Mini SAS HD connectors
- Attributes required: Available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECDV) 4.5M SAS YO12 Active Optical Cable (Adapter to Enclosure)

This 4.5 meter SAS cable connects one SAS adapter to a SAS I/O enclosure. This YO cable has three Mini-SAS HD (High Density) connectors. One of connectors attach to a SAS adapter such as in the PCIe3 SAS Adapter. The other two connectors attach to one SAS I/O Enclosure such as the EXP12SX (#ESLL/#ELLL) or EXP24SX (#ESLS/ #ELLS). This cable is designed for high speed (DHS) to support up to 12Gb throughput if the adapter has that capability.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/ O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters

can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: This cable cannot be used with the EXP24S I/O drawer (#5887 or# EL1S) which uses Mini-SAS connectors which are not HD.
Note: AOC cables require minimum level of SAS Adapter firmware. Refer to the latest prerequisites at

https://www14.software.ibm.com/webapp/set2/iprt/home

- Attributes provided: Connection between a SAS adapters with Mini-SAS HD connectors and a SAS I/O drawer with Mini SAS HD connectors
- Attributes required: Available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECDW) 10M SAS YO12 Active Optical Cable (Adapter to Enclosure)

This 10 meter SAS cable connects one SAS adapter to a SAS I/O enclosure. This YO cable has three Mini-SAS HD (High Density) connectors. One of connectors attach to a SAS adapter such as in the PCIe3 SAS Adapter. The other two connectors attach to one SAS I/O Enclosure such as the EXP12SX (#ESLL/#ELLL) or EXP24SX (#ESLS/ #ELLS). This cable is designed for high speed (DHS) to support up to 12Gb throughput if the adapter has that capability.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the adapter and the I/ O drawer. See SAS cabling documentation for the length of each leg of the cable. The SAS adapters can be in the same or in different PCIe I/O drawers. Or one adapter can be in a supported Power System CEC and the other adapter can be in a PCIe I/O drawer. Or both adapters can be in a supported Power System CEC.

Limitation: This cable cannot be used with the EXP24S I/O drawer (#5887 or# EL1S) which uses Mini-SAS connectors which are not HD.
Note: AOC cables require minimum level of SAS Adapter firmware. Refer to the latest prerequisites at

https://www14.software.ibm.com/webapp/set2/iprt/home

- Attributes provided: Connection between a SAS adapters with Mini-SAS HD connectors and a SAS I/O drawer with Mini SAS HD connectors
- Attributes required: Available connectors on SAS controllers and SAS I/O drawer
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECE0) 0.6M SAS AA12 Cable (Adapter to Adapter)

This 0.6 meter SAS cable connects a pair of SAS adapters to each other. This AA cable has two Mini-SAS HD (High Density) connectors which connect the top connectors of two PCIe3 SAS adapters with write cachce such as #EJ0L or #EJ14. The cable provides a high performance path of all the dual controller communication including mirroring the write cache and status awareness of each card. This cable is designed for high speed to support up to 12Gb throughput. Note EJ0L/EJ14 support 6Gb.

Two AA cables are always required between a pair of PCIe3 SAS adapters with write cache when just one or two I/O drawers are attached. One AA cable is required if three I/O drawers are attached. If four drawers are attached or if an AA cable fails or is disconnected, then the information it was carrying is then sent over the cables attached to the I/O drawers, sharing their bandwidth. AA cables are not used with SAS adapters with no write cache.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the

two adapters.

The 6Gb version of this cable is feature #5918. #5918 and #ECE0 can be mixed on the same PCIe3 adapter pair.

- Attributes provided: Connection between two SAS adapters with Mini-SAS HD connectors
- Attributes required: Available connectors on SAS controllers
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECE3) 3.0M SAS AA12 Cable

This 3.0 meter SAS cable has two Mini-SAS HD (High Density) connectors, and is designed for high speed to support up to 12Gb throughput. This is a straight cable (in contrast with X or YO cables) that has two distinct uses:

1. For Elastic Storage Server (ESS) solutions that have a 5147-024 I/O drawer, this cable is used to attach the 5147-024 to its controller.
2. For POWER Servers with #5887, #EL1S, #ESLS, #ESLL, #ELLS, or #ELLL I/O drawers driven by paired PCIe controllers with write cache such as #EJ0L or #EJ14, this cable is used to connect the top connectors of the paired controllers. The cable provides a high performance path of all the dual controller communication including mirroring the write cache and status awareness of each card.

Note that X or YO cables are always used to attach I/O drawers #5887, #EL1S, #ESLS, #ESLL, #ELLS, or #ELLL to controllers on POWER Servers. Straight cables (such as #ECE3) are not allowed to directly attach to I/O drawers on POWER Servers.

Two AA cables are always required between a pair of PCIe3 SAS adapters with write cache when just one or two I/O drawers are attached. One AA cable is required if three I/O drawers are attached. If four drawers are attached or if an AA cable fails or is disconnected, then the information it was carrying is then sent over the cables attached to the I/O drawers, sharing their bandwidth. On POWER Servers, AA cables are not used with SAS adapters with no write cache.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the two adapters.

The 6Gb version of this cable is feature #5915. #5915 and #ECE3 can be mixed on the same PCIe3 adapter pair.

- Attributes provided: For ESS solutions, connection between a SAS controller and one 5174-024 I/O drawer. For POWER Systems, connection between two paired SAS controllers with write cache and Mini-SAS HD connectors.
- Attributes required: For ESS solutions, a 5147-024 I/O drawer and appropriate controller. For POWER Systems, available connectors on SAS controllers.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECE4) 4.5M SAS AA12 Active Optical Cable (Adapter to Adapter)

This 4.5 meter SAS cable connects a pair of SAS adapters to each other. This AA cable has two Mini-SAS HD (High Density) connectors which connect the top connectors of two PCIe3 SAS adapters with write cachce such as #EJ0L or #EJ14. The cable provides a high performance path of all the dual controller communication including mirroring the write cache and status awareness of each card. This cable is designed for high speed to support up to 12Gb throughput. Note EJ0L/EJ14 support 6Gb.

Two AA cables are always required between a pair of PCIe3 SAS adapters with write cache when just one or two I/O drawers are attached. One AA cable is required if three I/O drawers are attached. If four drawers are attached or if an AA cable fails or is disconnected, then the information it was carrying is then sent over the cables attached to the I/O drawers, sharing their bandwidth. AA cables are not used with SAS adapters with no write cache.

Multiple cable length feature codes are available. Choose the cable length that best matches the distance between the two adapters.
Note: AOC cables require minimum level of SAS Adapter firmware. Refer to the latest prerequisites at

https://www14.software.ibm.com/webapp/set2/iprt/home

- Attributes provided: Connection between two SAS adapters with Mini-SAS HD connectors
- Attributes required: Available connectors on SAS controllers
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECJ5) 4.3m (14-Ft) PDU to Wall 3PH/24A 200-240V Delta-wired Power Cord

This power cord feature ECJ5 contains an amphenol type of connector and only supported on PDUs ECJK or ECJL, and ECJP or ECJQ.

ECJ5 has a 4-pin IEC 60309 style plug, 430P9W. It contains three line conductors and a protective earth, but no neutral. ECJ5 is supported in countries that use a delta electrical distribution. ECJ5 is not supported in China, Hong Kong, and other countries that use a wye electrical distribution.

- Attributes provided: Power cord
- Attributes required: PDU features ECJK or ECJL, and ECJP or ECJQ.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECJ6) 4.3m (14-Ft) PDU to Wall 3PH/40A 200-240V Power Cord

This power cord goes from the chassis to a wall-type outlet. The line cord has a 200-240V, CS8365 50A Plug (de-rated to 40A) and Amphenol inlet compatible with PDU FCs ECJL and ECJQ. The following countries/regions use the #ECJ6 power cord to power the system: United States and Canada only.

- Attributes provided: Power cord PDU to Wall
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECJ7) 4.3m (14-Ft) PDU to Wall 3PH/48A 200-240V Delta-wired Power Cord

This power cord feature ECJ7 contains an amphenol type of connector and only supported on PDUs ECJK or ECJL, and

ECJP or ECJQ.

ECJ7 has a 4-pin IEC 60309 style plug, 460P9W. It contains three line conductors and a protective earth, but no neutral. ECJ7 is supported in countries that use a delta electrical distribution. ECJ7 is not supported in China, Hong Kong, and other countries that use a wye electrical distribution.

- Attributes provided: Power cord
- Attributes required: PDU features ECJK or ECJL, and ECJP or ECJQ.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECJJ) High Function 9xC19 Single-Phase or Three-Phase Wye PDU plus

This is an intelligent, switched 200-240 volt single-phase or 380-415/220-240 volt three-phase wye AC Power Distribution Unit (PDU) plus with nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the nine C19 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Depending on country wiring standards the PDU is single-phase or three-phase wye. Three-phase wye-wired connectors have 5-pins and use three line conductors, a neutral, and a protective earth. The input is 380-415 volt line-to-line and the output is 220-240 volt line-to-neutral for three-phase wye PDUs.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C20 plug connect to C19 PDU receptacles and are ordered separately. One country-specific wall line cord is also ordered separately and attaches to a UTG524-7 connector on the front of the PDU. Supported line cords include features #6489, #6491, #6492, #6653, #6654, #6655, #6656, #6657, #6658, #6667, #ELC1 or #ELC2.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

There are also three C13 receptacles on the rear of the PDU positioned toward the middle of the rack. These are generally not easily accessed and therefore IBM does not generally recommend their use.

#ECJG and #ECJJ are identical PDUs. Up to one lower price #ECJG can be ordered with a new 7965-S42 rack in place of a no-charge #9188 PDU.

For comparison, this is most similar to the earlier generation #EPTJ PDU.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Nine C19 PDU plus - switched, power monitoring
- Attributes required: PDU wall line cord and space in 19-inch rack
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

1h5.(#ECJL) - High Function 9xC19 PDU plus 3-Phase Delta

This is an intelligent, switched 200-240 volt 3-phase delta AC Power Distribution Unit (PDU) plus with nine C19

receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the nine C19 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Three-phase delta-wired connectors have 4-pins and use three line conductors and a protective earth. The input is 200-240 volt line-to-line and the output is 200-240 volt line-to-line for three-phase delta PDUs.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C20 plug connect to C19 PDU receptacles and are ordered separately. One wall line cord is also ordered separately and attaches to the Amphenol inlet connector. Supported line cords include features #ECJ5, #ECJ6, and #ECJ7.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

There are also three C13 receptacles on the rear of the PDU positioned toward the middle of the rack. These are generally not easily accessed and therefore IBM does not generally recommend their use.

#ECJK and #ECJL are identical PDUs. Up to one lower price #ECJK can be ordered with a new 7014-T42/T00 rack in place of a no-charge #9188 PDU.

For comparison, this is most similar to the earlier generation #EPTL PDU.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Nine C19 PDU plus - switched, power monitoring
- Attributes required: space in rack, 3-phase 208V AC delta electrical service
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECJN) High Function 12xC13 Single-Phase or Three-Phase Wye PDU plus

This is an intelligent, switched 200-240 volt single-phase or 380-415/220-240 volt three-phase wye AC Power Distribution Unit (PDU) plus with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the twelve C13 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Depending on country wiring standards the PDU is single-phase or three-phase wye. Three-phase wye-wired connectors have 5-pins and use three line conductors, a neutral, and a protective earth. The input is 380-415 volt line-to-line and the output is 220-240 volt line-to- neutral for three-phase wye PDUs.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C14 plug connect to C13 PDU receptacles and are ordered separately. One country-specific wall line cord is also ordered separately and attaches to a UTG524-7 connector on the front of the PDU. Supported line cords include features #6489, #6491, #6492, #6653, #6654, #6655, #6656, #6657, #6658, #6667, #ELC1 or #ELC2.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

Feature #ECJM and #ECJN are identical PDUs. Up to one lower price #ECJM can be ordered with a new 7014-T42/T00 rack in place of a no- charge #9188 PDU.

For comparison, this is most similar to the earlier generation #EPTN PDU.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Twelve C13 PDU plus - switched, power monitoring
- Attributes required: PDU wall line cord and space in 19-inch rack
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECJQ) High Function 12xC13 PDU plus 3-Phase Delta

This is an intelligent, switched 200-240 volt 3-phase delta AC Power Distribution Unit (PDU) plus with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the twelve C13 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Three-phase delta-wired connectors have 4-pins and use three line conductors and a protective earth. The input is 200-240 volt line-to-line and the output is 200-240 volt line-to-line for three-phase delta PDUs.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C20 plug connect to C19 PDU receptacles and are ordered separately. One wall line cord is also ordered separately and attaches to the Amphenol inlet connector. Supported line cords include features #ECJ5, #ECJ6 and #ECJ7.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

#ECJP and #ECJQ are identical PDUs. Up to one lower price #ECJP can be ordered with a new 7965-S42 rack in place of a no-charge #9188 PDU.

For comparison, this is most similar to the earlier generation #EPTP PDU.

Not supported in China, Hong Kong, and other countries that use a wye electrical distribution.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Twelve C13 PDU plus - switched, power monitoring
- Attributes required: space in rack, 3-phase 208V AC delta electrical service
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ECSF) Custom Service Specify, Montpellier, France

Having #ECSF on the order, will cause the order to be routed to France and the machine to be internally routed to the CSC build area.

- Attributes provided: Customization
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ECSM) Custom Service Specify, Mexico

Having #ECSM on the order, will cause the order to be routed to Mexico and the machine to be internally routed to the CSC build area.

- Attributes provided: Customization
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: N/A
- Return parts MES: Does not apply

(#ECSP) Custom Service Specify, Poughkeepsie, USA

Having #ECSP on the order, will cause the order to be routed to Poughkeepsie, USA and the machine to be internally routed to the CSC build area.

- Attributes provided: Customization
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: N/A
- Return parts MES: Does not apply

(#ECSS) Integrated Solution Packing

This is a routing indicator for Solution packing.

- Attributes provided: Routing instruction for manufacturing
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: N/A
- Return parts MES: Does not apply

(#ECW0) Optical Wrap Plug

A wrap plug is a small connector designed to perform a diagnostic test called a loopback test. This wrap plug is inserted into a SR optical port on a PCIe Fibre Channel adapter or a SR or LR optical port on a PCIe Ethernet adapter

This is a multi-mode LC fiber optic wrap plug with an inside/ outside optics diameter of 50/125. Its IBM part number as of early 2016 is 12R9314. An earlier equivalent function IBM part number which is no longer shipped is 11P3847.

It is strongly recommended that Fibre Channel adapters (HBAs) fill any empty adapter ports with a wrap plug. There is no technical issue leaving a port empty. However, filling all ports with a cable to a device/switch or with a wrap plug can speed the booting/IPLing of a partition and can avoid error messages uselessly pointing to a planned empty port.

There is no technical issue leaving an Ethernet port empty. Whether an Ethernet port is empty or contains a wrap plug should not impact boot/IPL time or impact empty-port messages.

- Attributes provided: Wrap plug
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ED2Z) Single 5250 Enterprise Enablement

Feature provides one processor cores worth of 5250 OLTP capacity which can be spread across multiple physical processor cores or multiple partitions. A permanently activated processor core and adequate IBM i processor license entitlements are prerequisites.

- Attributes provided: One Processor of 5250 Enterprise Enablement
- Attributes required: IBM i environment
- Minimum required: 0
- Maximum allowed: 64 (Initial order maximum: 64)
- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ED30) Full 5250 Enterprise Enablement

Feature provides 5250 OLTP capacity for all permanently activated processor cores. Adequate IBM i processor licenses are a prerequisite for the 5250 OLTP usage.

- Attributes provided: Full 5250 Enterprise Enablement
- Attributes required: IBM i environment
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EDAB) 100 GB DDR4 Mobile Memory Activation for HEX

One #EDAB feature is ordered for each 100 GB DDR4 mobile memory activation.

- Attributes provided: Enabled activation of 100 GB DDR4 mobile memory.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 480 (Initial order maximum: 480)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EDAG) 256 GB Base Memory Activation (Pools 2.0)

Feature EDAG allows the activation of 256 GB Memory for Pools 2.0

- Attributes provided: 256 Memory Activation for Pools 2.0
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 256 (Initial order maximum: 256)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAG not orderable in China

(#EDAH) 512 GB Base Memory Activation (Pools 2.0)

Feature EDAH allows the activation of 512 GB Memory for Pools 2.0

- Attributes provided: 512 Memory Activation for Pools 2.0
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAH not orderable in China

(#EDAL) 256 GB Base Memory Activation Linux only

Feature EDAL allows the activation of 256 GB Memory for Linux

- Attributes provided: 256 Memory Activation for Linux only
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 256 (Initial order maximum: 256)
- OS level required:
  - IBM i - Not Supported
  - AIX - Not Supported
- Initial Order/MES/Both/Supported: Both

- CSU: Yes
- Return parts MES: No
  Note: Feature EDAL not orderable in China

(#EDAM) 512 GB Base Memory Activation Linux only

Feature EDAM allows the activation of 512 GB Memory for Linux

- Attributes provided: 512 Memory Activation for Linux only
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required:
  - IBM i - Not Supported
  - AIX - Not Supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAM not orderable in China

(#EDAP) 1 GB Base Memory activation (Pools 2.0) from Static

This feature delivers 1 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 1 GB Base Memory Activation for Pools 2.0.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 65536 (Initial order maximum: 65536)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAP not orderable in China

(#EDAQ) 100 GB Base Memory activation (Pools 2.0) from Static

This feature delivers 100 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 100 GB Base Memory Activation for Pools 2.0.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 655 (Initial order maximum: 655)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAQ not orderable in China

(#EDAR) 512 GB Base Memory activation (Pools 2.0) from Static

This feature delivers 512 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 512 GB Base Memory Activation for Pools 2.0.

-
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAR not orderable in China

(#EDAS) 500 GB Base Memory activation (Pools 2.0) from Static

This feature delivers 500 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 500 GB Base Memory Activation for Pools 2.0.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 130 (Initial order maximum: 130)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAS not orderable in China

(#EDAT) 1 GB Base Memory activation (Pools 2.0) MES only

This feature delivers 1 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 1 GB Base Memory Activation for Pools 2.0.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 65536 (Initial order maximum: 65536)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAT not orderable in China

(#EDAU) 100 GB Base Memory activation (Pools 2.0) MES only

This feature delivers 100 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 100 GB Base Memory Activation for Pools 2.0.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 655 (Initial order maximum: 655)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAU not orderable in China

(#EDAV) 100 GB Base Memory Activation (Pools 2.0) from Mobile

This feature delivers 100 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 100 GB Base Memory Activation for Pools 2.0.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 655 (Initial order maximum: 655)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAV not orderable in China

(#EDAW) 500 GB Base Memory Activation (Pools 2.0) from Mobile

This feature delivers 500 GB Base Memory only for Pools 2.0 on HEX

- Attributes provided: 500 GB Base Memory Activation for Pools 2.0.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 130 (Initial order maximum: 130)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAW not orderable in China

(#EDAX) 512 GB Base Memory Activation Linux only - Conversion

Feature EDAX allows the activation of 512 GB Memory for Linux

- Attributes provided: 512 Memory Activation for Linux only
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required:
  - IBM i - Not Supported
  - AIX - Not Supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDAX not orderable in China

(#EDBK) 2 x SMP cable brackets for non-IBM Rack

- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: No

- Return parts MES: No

(#EDFP) Flexible service processor

- Minimum required: 0
- Maximum allowed: 2 (Initial order maximum: 2)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EDN1) 5U System node Indicator drawer

This feature provides a 5U system node drawer. It includes the drawer, labels, 4 CPU Sockets, 64 DDIMM Slots, 8 Low profile PCIe Slots, 4 NVME slots.

- Attributes provided: One 5U system node drawer
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EDP2) 40-core (4x10) 3.65 to 3.90 GHz (max) Power10 Processor with 5U system node drawer

Fourty core 3.65 to 3.90 GHz (max) Power10 processor planar containing four ten-core processor SCMs. Each processor core has 2 MB of L2 cache and each SCM has a total of 80MB of L3 cache. There are 16 DDIMM slots per processor and a total of 64 DDIMM slots per system node. Each processor feature will deliver a set of four identical processors that populate all the sockets in a given system node. All processors in the system node must be identical. Activations are applied per core with a minimum of sixteen activations.

- Attributes provided: Four 10-core 3.65 to 3.90x GHz (max) SCMs providing 40-core system node
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
  OS level required:
  - IBM i 7.5, or later
  - IBM i 7.3 TR11 or later Supported
  - IBM i 7.4 TR5 or later Supported
  - AIX - Supported
  - Linux - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EDP3) 48-core (4x12) 3.60 to 4.15 GHz (max) Power10 Processor with 5U system node drawer

Fourty-eight core 3.60 to 4.15 GHz (max) Power10 processor planar containing four twelve-core processor SCMs. Each processor core has 2 MB of L2 cache and each SCM has a total of 120MB of L3 cache. There are 16 DDIMM slots per processor and a total of 64 DDIMM slots per system node. Each processor feature will deliver a set of four identical

processors that populate all the sockets in a given system node. All processors in the system node must be identical. Activations are applied per core with a minimum of sixteen activations.

- Attributes provided: Four 12-core 3.60 to 4.15 GHz (max) SCMs providing 48-core system node
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
  OS level required:
  - IBM i 7.5, or later
  - IBM i 7.3 TR11 or later Supported
  - IBM i 7.4 TR5 or later Supported
  - AIX - Supported
  - Linux - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EDP4) 60-core (4x15) 3.55 to 4.00 GHz (max) Power10 Processor with 5U system node drawer

Sixty core 3.55 to 4.00 GHz (max) Power10 processor planar containing four fifhteen-core processor SCMs. Each processor core has 2 MB of L2 cache and each SCM has a total of 120MB of L3 cache. There are 16 DDIMM slots per processor and a total of 64 DDIMM slots per system node. Each processor feature will deliver a set of four identical processors that populate all the sockets in a given system node. All processors in the system node must be identical. Activations are applied per core with a minimum of sixteen activations.

- Attributes provided: Four 15-core 3.55 to 4.00 GHz (max) SCMs providing 60-core system node
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
  OS level required:
  - IBM i 7.5, or later
  - IBM i 7.3 TR11 or later Supported
  - IBM i 7.4 TR5 or later Supported
  - AIX - Supported
  - Linux - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No
  Note: Feature EDP4 not orderable in China

(#EDPB) 1 core Processor Activation for #EDP2

Each occurrence of this feature will permanently activate one processor core on Processor Card #EDP2

- Attributes provided: One processor core activation for #EDP2
- Attributes required: #EDP2 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 160 (Initial order maximum: 160)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes

- Return parts MES: No

(#EDPC) 1 core Processor Activation for #EDP3

Each occurrence of this feature will permanently activate one processor core on Processor Card #EDP3

- Attributes provided: One processor core activation for #EDP3
- Attributes required: #EDP3 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EDPD) 1 core Processor Activation for #EDP4

Each occurrence of this feature will permanently activate one processor core on Processor Card #EDP4

- Attributes provided: One processor core activation for #EDP4
- Attributes required: #EDP4 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EDPD not orderable in China

(#EDPZ) Mobile processor activation for HEX

Mobile processor activation upgrade for HEX

- Attributes provided: Processor activation code number
- Attributes required: Processor FCs
- Minimum required: 0
- Maximum allowed: 224 (Initial order maximum: 224)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EFCE) System Node to System Control Unit Cable Set for Drawer 2

Provides cables required to connect a second system node with the system control unit. Includes: 2 x 1m FSP Cable (Right & Left), 4 x 0.88m SMP Cable, 8 x 1.005m SMP Cable, 4 x 1.225m SMP Cable (there are a total of 2 UPIC cables that are inlcuded in EFCA).

- Attributes provided: System Node to System Control Unit Cable set for Drawer 2
- Attributes required: System Node Drawer 2 and System Control Unit
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)

- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EFCF) System Node to System Control Unit Cable Set for Drawer 3

Provides cables required to connect a third system node with the system control unit. Includes: 2 x 1m FSP Cable (Right & Left), 8 x 1.225m SMP Cable.

- Attributes provided: System Node to System Control Unit Cable set for Drawer 3
- Attributes required: System Node Drawer 3 and System Control Unit
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EFCG) System Node to System Control Unit Cable Set for Drawer 4

Provides cables required to connect a fourth system node with the system control unit. Includes: 2 x 1m FSP Cable (Right & Left), 4 x 0.880 SMP Cable, 4 x 1.005m SMP Cable, 8 x 1.225m SMP Cable, 4 x 1.425m SMP Cable, 4 x 1.625m SMP Cable.

- Attributes provided: System Node to System Control Unit Cable set for Drawer 4
- Attributes required: System Node Drawer 4 and System Control Unit
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EFCH) System Node to System Control Unit Cable Set for Drawer 1

Provides cables required to connect the system node with the system control unit; this configuration does not require External SMP cables. Includes: 2 x 1m FSP Cable (Right & Left), 2 x 330mm Power/UPIC Cable (Right & Left), 1.05m USB Cable.

- Attributes provided: System Node to System Control Unit cable set for Drawer1.
- Attributes required: System Node Drawer 1 and System Control Unit
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EFCR) Captive Rack identifier

- Minimum required: 0

- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EH35) Mobile Enablement

Enterprise Pool Enablement

- Attributes provided: Enterprise Pool Enablement Code
- Attributes required: none
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 0)
- OS level required:
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: Does not apply

(#EHKV) SAP HANA TRACKING FEATURE

SAP HANA tracking feature that defines manufacturing routing.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required:
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EHR2) Boot Drive / Load Source in EXP24SX Specify (in #ESLS or #ELLS)

Indicates that boot drive or load source (disks or SSDs)are placed in an EXP24SX SAS Storage Enclosure

- Attributes provided: Boot drive / load source location specify
- Attributes required: Available SAS bay and supported disk/SSD
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EHS2) SSD Placement Indicator - #ESLS/#ELLS

This is an IBM internal automatic generated SSD specify indicator for placement and it is not selectable.

- Attributes provided: None
- Attributes required: None

- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required:
  - AIX - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJ0J) PCIe3 RAID SAS Adapter Quad-port 6Gb x8

The PCIe3 RAID SAS Adapter is a high performance SSD/HDD controller using PCIe Gen3 x8 technology. The adapter does not have write cache and thus pairing with another PCIe3 RAID SAS Adapter (#EJ0J or #EJ0M) is optional. Pairing can provide controller redundancy and enhance performance. There are no batteries in the adapter to maintain.

The adapter provides four Mini-SAS HD (high density) narrow connectors for the attachment of SAS drives located in the EXP24S, EXP12SX, or EXP24SX storage enclosures or #5802/5803/EL36 12X PCIe I/ O drawers. X, YO or AT SAS cables with HD narrow connectors are used to attach to these drawers. A max of 4 EXP24S/EXP12SX/EXP24SX can be attached. A maximum of 48 SSD can be attached and a maximum of 96 HDD can be attached per adapter or per adapter pair.

The adapter provides RAID 0, RAID 5, RAID 6 and RAID 10 for AIX and Linux and VIOS. The adapter provides RAID 5 and RAID 6 for all levels of IBM i and also provides RAID 10 for later levels of IBM i. IBM i provides both OS mirroring and data spreading. AIX/Linux/VIOS provide OS mirroring (LVM).

Features #EJ0J and #EJ0M are electronically identical with the same CCIN of 57B4. #EJ0J has a full-high tailstock bracket and air baffle. #EJ0M has a low profile tailstock bracket. #EJ10/#EJ11 are identical with #EJ0J/#EJ0M, but have different feature codes to identify their use as tape/DVD controllers to IBM configurator tools instead of disk/SSD controllers.

Both 5xx and 4k byte sector HDD/SSD are supported for POWER8/POWER9 servers. 5xx byte sector HDD/SSD are supported for earlier generation servers. 5xx and 4k drives cannot be mixed in the same array.

Limitations:

- HDD/SSD workloads which are performance sensitive to WRITES should use the #EJ14 or #EJ0L controller which provides write cache.
- HDD and SSD cannot be mixed on the same SAS port, but can be mixed on the same adapter.
- #5886 3.5-inch SAS Storage Drawer is not supported.
- 177GB SSD are not supported.
- Running SAS bays for both a #5887 EXP24S I/O drawer and a 12X-attached #5802/5803 I/O drawer on the same adapter or adapter pair is not supported. Note mixing EXP24S or EXP12SX or EXP24SX is supported.
- If controlling drives in a #5802/5803/EL36 as a single controller, the #EJ0J must be located in that #5802/5803/EL36. If controlling drives in a #5802/5803/EL36 as a pair of controllers, at least one of the SAS adapter pairs must be located in that #5802/ 5803/EL36.
- Tape/DVD cannot be mixed with disk/SSD on the same adapter.
- Attributes provided: full high PCIe3 four port x8 SAS RAID adapter with no write cache and optional pairing
- Attributes required: One PCIe slot per adapter and Mini-SAS HD narrow connector SAS cables such as #ECBJ-ECBL, #ECBT-ECBV, #ECCO-ECC4 or 12Gb cables such as #ECDJ, ECDT or #ECDU.
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both

- CSU: Yes
- Return parts MES: No
  - VIOS supported

(#EJ0M) PCIe3 LP RAID SAS Adapter Quad-Port 6Gb x8

The PCIe3 RAID SAS LP Adapter is a high performance SSD/HDD controller using PCIe Gen3 x8 technology. The adapter does not have write cache and thus pairing with another PCIe3 RAID SAS Adapter (#EJ0M or #EJ0J) is optional. Pairing can provide controller redundancy and enhance performance. There are no batteries in the adapter to maintain.

The adapter provides four Mini-SAS HD (high density) narrow connectors for the attachment of SAS drives located in the EXP24S, EXP12SX, or EXP24SX storage enclosure or #5802/EL36 12X PCIe I/O drawers. X, YO or AT SAS cables with HD narrow connectors are used to attach to these drawers. A max of 4 EXP24S/EXP12SX/EXP24SX can be attached. A maximum of 48 SSD can be attached and a maximum of 96 HDD can be attached per adapter or per adapter pair.

The adapter provides RAID 0, RAID 5, RAID 6 and RAID 10 for AIX and Linux and VIOS. The adapter provides RAID 5 and RAID 6 for all levels of IBM i and also provides RAID 10 for later levels of IBM i. IBM i provides both OS mirroring and data spreading. AIX/ Linux/VIOS provide OS mirroring (LVM).

Features #EJ0J and #EJ0M are electronically identical with the same CCIN of 57B4. #EJ0J has a full-high tailstock bracket and air baffle. #EJ0M has a low profile tailstock bracket. #EJ10 / #EJ11 are identical with #EJ0J / #EJ0M respectively, but have different feature codes to identify their use to IBM configurator tools as tape/DVD controllers instead of disk/SSD controllers.

Both 5xx and 4k byte sector HDD/SSD are supported for POWER8/POWER9 servers. 5xx byte sector HDD/SSD are supported for earlier generation servers. 5xx and 4k drives cannot be mixed in the same array.

Limitations:

- HDD/SSD workloads which are performance sensitive to WRITES should use the #EJ14 or #EJ0L controller which provides write cache.
- HDD and SSD cannot be mixed on the same SAS port, but can be mixed on the same adapter.
- #5886 3.5-inch SAS Storage Drawer is not supported.
- 177GB SSD are not supported.
- Running SAS bays for both a EXP24S storage drawer and a 12X-attached #5802/5803 I/O drawer on the same adapter or adapter pair is not supported. Note mixing EXP24S or EXP12SX or EXP24SX is supported.
- A single #EJ0M cannot control drives in a #5802/5803/EL36 because the adapter must be located in that #5802/5803/ EL36 in a full high slot. However, if controlling drives in a #5802/5803/EL36 as a pair of controllers, one #EJ0J can be located in that #5802/ 5803/EL36 and one #EJ0M can be in the system's low profile slot.
- Tape/DVD cannot be mixed with disk/SSD on the same adapter.
- Attributes provided: low profile PCIe3 four port x8 SAS RAID adapter with no write cache and optional pairing
- Attributes required: One low profile PCIe slot per adapter and Mini-SAS HD narrow connector SAS cables such as #ECBJ-ECBL, #ECBT-ECBV, #ECCO-ECC4 or 12Gb cables such as #ECDJ, ECDT or #ECDU.
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  - VIOS supported

(#EJ10) PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8

The PCIe3 SAS Adapter is a high performance SAS tape controller using PCIe Gen3 x8 technology. The adapter supports external SAS tape drives such as the LTO-5, LTO-6, LTO-7, and LTO-8 found in the IBM 7226-1U3 Multimedia drawers, or tape units such as the TS2250, TS2260, TS2270, and TS2280 single External Tape Drive, TS2900, TS3100, TS3200, and TS3310. Other removable media devices supported include IBM SAS/SATA DVD RAM drive features available on the IBM 7226-1U3 Storage Enclosure. The adapter provides four Mini-SAS HD (high density) connectors to which AE1 SAS cables such as #ECBY and/ or YE1 SAS Cables such as #ECBZ with HD narrow connectors can be attached. A max of 4 tape drives per adapter can be attached using four AE1 cables. A max of 8 tape drives can be attached using four YE1 cables.

#EJ10 (full high) and #EJ11 (low profile) are electronically the same adapter with the same 57B4 CCIN, but differ in that their tailstocks fit different size PCIe slots.

#EJ0J and #EJ10 are the same adapter with the same 57B4 CCIN, but have different feature code numbers to indicate different usage to IBM configurator tools. #EJ10 runs SAS LTO-5 or later tape drives and DVD. Support of both tape/DVD and HDD/SSD on the same adapter is not supported.
Note: The original #EJ0X adapter does not support DVD but also has the same CCIN.
Note: Adapter uses a Mini-SAS HD narrow connector and AE1 #ECBZ or YE1 #ECBY SAS cable.

Limitation: LTO-4 or earlier drives are not supported.

- Attributes provided: full high PCIe3 four port x8 SAS adapter
- Attributes required: One PCIe slot per adapter
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  - VIOS supported

(#EJ11) PCIe3 LP SAS Tape/DVD Adapter Quad-port 6Gb x8

The PCIe3 Low Profile SAS Adapter is a high performance SAS tape controller using PCIe Gen3 x8 technology. The adapter supports external SAS tape drives such as the LTO-5, LTO-6, LTO-7, and LTO-8 found in the IBM 7226-1U3 Multimedia drawers, or tape units such as the TS2250, TS2260, TS2270, and TS2280 single External Tape Drive, TS2900, TS3100, TS3200, and TS3310. Other removable media devices supported include IBM SAS/SATA DVD RAM drive features available on the IBM 7226-1U3 Storage Enclosure. The adapter provides four Mini-SAS HD (high density) connectors to which AE1 SAS cables such as #ECBY and/or YE1 SAS Cables such as #ECBZ with HD narrow connectors can be attached. A max of 4 tape drives per adapter can be attached using four AE1 cables. A max of 8 tape drives can be attached using four YE1 cables.

#EJ10 (full high) and #EJ11 (low profile) are electronically the same adapter with the same 57B4 CCIN, but differ in that their tailstocks fit different size PCIe slots.

#EJ0M and #EJ11 are the same adapter with the same 57B4 CCIN, but have different feature code numbers to indicate different usage to IBM configurator tools. #EJ11 runs SAS LTO-5 or later tape drives and DVD. Support of both tape/DVD and HDD/SSD on the same adapter is not supported.
Note: Adapter uses a Mini-SAS HD narrow connector and AE1 #ECBZ or YE1 #ECBY SAS cable.

Limitation: LTO-4 or earlier drives are not supported.

- Attributes provided: low profile PCIe3 four port x8 SAS adapter
- Attributes required: One low profile PCIe slot per adapter
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)

  OS level required:

- AIX - supported
- IBM i - supported
- Linux - supported

Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  - VIOS supported

(#EJ14) PCIe3 12GB Cache RAID PLUS SAS Adapter Quad-port 6Gb x8

The PCIe3 12GB Cache RAID PLUS SAS Adapter provides high performance HDD and/or SSD controller function using PCIe Gen3 technology. A pair of #EJ14 adapters are required to provide mirrored write cache data and adapter redundancy. Integrated flash memory provides protection of the write cache without batteries in case of power failure. Effectively up to 12GB of write cache is provided using compression of 4 GB of physical cache. The #EJ14 provides four Mini-SAS HD (high density) narrow connectors for the attachment of SAS drives located in the SAS EXP24S, EXP12SX, or EXP24SX storage enclosures. X, YO or AT SAS cables with HD narrow connectors are used to attach to these drawers. A max of 4 EXP24S can be attached with a maximum of 96 HDD or a maximum of 72 SSD per pair of #EJ14. If more than 48 SSD are attached, no HDD can be attached. A mix of EXP24S, EXP12SX or EXP24SX is supported on the same adapter pair. Two AA SAS cable with HD narrow connectors are attached to the #EJ14 pair to communicate status and cache content information and are required unless three or four ports are being used to attach HDD/SSD. Feature #EJ14 provides RAID 0, RAID 5, RAID 6, and RAID 10, RAID 5T2, RAID 6T2, and RAID 10T2 for AIX and Linux and VIOS. Two tier arrays (5T2, 6T2 and 10T2) combine both HDD and SSD into a single array with Easy Tier functionality. AIX/Linux/VIOScan also provide OS mirroring (LVM). On systems that support IBM i, the adapter provides RAID 5 and RAID 6 for IBM i. RAID 10 is supported by IBM i 7.2. IBM i provides both OS mirroring and data spreading. This adapter is very similar to the#EJ0L SAS adapter, but #EJ14 uses a second CPU chip in the card to provide more IOPS capacity and can attach more SSD. The #EJ14 adapter's CCIN is 57B1. Both 5xx and 4k byte sector HDD/SSD are supported. 5xx and 4k drives cannot be mixed in the same array. Limitations:

- Not supported on POWER7/POWER7+ servers.
- HDD and SSD cannot be mixed on the same SAS port, but can be mixed on the same adapter.
- Attributes provided: Full high PCIe3 four port x8 adapter with up to 12 GB write cache
- Attributes required: One PCIe slot per adapter and Mini-SAS HD narrow connector SAS cables such as #ECBJ-ECBL, #ECBT-ECBV, #ECCO-ECC4 or 12Gb cables such as ECDT or #ECDU.
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: VIOS supported

(#EJ24) PCIe x16 to CXP Converter Card, Supports optical cables

PCIe3 x16 adapter provides two optical CXP ports for the attachment of two active optical cables (AOC). One adapter supports the attachment of one PCIe3 module in a PCIe Gen3 I/O Expansion Drawer. CCIN is 6B53.

- Attributes provided: PCIe3 adapter with two CXP ports to attach two active optical cables
- Attributes required: PCIe3 x16 slot in system unit plus a pair of active optical cables (one optical pair feature such as #ECCR, #ECCY, or #ECCZ)

- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJ33) PCIe3 Crypto Coprocessor BSC-Gen3 4767

Secure-key adapter provides both cryptographic coprocessor and cryptographic accelerator functions in a single PCIe card. The adapter is well suited to applications requiring high-speed, security- sensitive, RSA acceleration, cryptographic operations for data encryption and digital signing, secure management, and use of cryptographic keys, or custom cryptographic applications. It provides secure storage of cryptographic keys in a tamper-resistant hardware security module designed to meet FIPS 140-2 level 4 security requirements. The adapter is a PCIe Gen 3 x4 full height - short card. The adapter runs in dedicated mode only (no PowerVM virtualization).

#EJ32 and #EJ33 are both feature codes representing the same physical card with the same CCIN of 4767. Different feature codes are used to indicate if a blind swap cassette is used and its type. #EJ32 indicates no blind swap cassette. #EJ33 indicates a Gen 3 blind swap cassette.

IBM PCIe Cryptographic Coprocessor adapter highlights

- Integrated Dual processors that operate in parallel for higher reliability, often 2x performance improvement over prior generation crypto cards
- Uses newer level Power Processor (PPC) processor than previous generation cards
- Supports IBM Common Cryptographic Architecture (CCA 5.3) and PKCS#11 standand
- Ability to configure adapter as coprocessor or accelerator
- Support for smart card applications using Europay, MasterCard and Visa
- Cryptographic key generation and random number generation
- PIN processing - generation, verification, translation
- Encrypt/Dectrypt using AES and DES keys

Please refer to the following URL for the latest firmware and software updates http://www-03.ibm.com/security/cryptocards/

- Attributes provided: Cryptographic Coprocessor and Accelerator Functions
- Attributes required: One full-high PCIe Gen3 slot which uses a blind swap cassette
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 0)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Linux software support can be downloaded from the following location: http://www-03.ibm.com/security/cryptocards/pciecc2/ ordersoftware.shtml

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EJ37) PCIe3 Crypto Coprocessor BSC-Gen3 4769

The 4769 PCIe Cryptographic Coprocessor has a PCIe local-bus- compatible interface. The coprocessor holds a security-enabled subsystem module and batteries for backup power. The hardened encapsulated subsystem contains two sets of

two 32-bit PowerPC 476FP reduced instruction set computer (RISC) processors running in lockstep with cross-checking to detect soft errors in the hardware. It also contains a separate service processor used to manage self-test and firmware updates; RAM; flash memory and battery-powered memory; secure time-of-day; cryptographic quality random number generator; AES; DES; Triple DES; HMAC; CMAC; MD5; multiple SHA hashing methods; modular- exponentiation hardware, such as RSA and ECC; and full-duplex direct memory access (DMA) communications.

A security-enabled code-loading arrangement allows control program and application program loading and refreshes after coprocessor installation in your server. IBM offers an embedded subsystem control program and a cryptographic application programming interface (API) that implements the IBM CCA.

The IBM Common Cryptographic Architecture Support Program can be accessed from the internet at no charge to the user. See the IBM CCA Basic Services Reference and Guide, which can be found at the IBM Cryptocards Library for a full explanation of the CCA API.

For details on future updates to the versions of operating systems that are supported by the 4769 PCIe Cryptographic Coprocessor, see the IBM Cryptocards website https://www.ibm.com/security/cryptocards.

Feature #EJ35 and #EJ37 are both feature codes representing the same physical card with the same CCIN of C0AF. Different feature codes are used to indicate if a blind swap cassette is used and its type. #EJ35 indicates no blind swap cassette. #EJ37 indicates a Gen 3 blind swap cassette.

The 4769 PCIe Cryptographic Coprocessor is designed to deliver the following functions:

- X.509 certificate services support
- ANSI X9 TR34-2019 key exchange services that exploit the public key infrastructure (PKI)
- ECDSA secp256k1
- CRYSTALS-Dilithium, a quantum-safe algorithm for digital signature generation and verification
- Rivest-Shamir-Adleman (RSA) algorithm for digital signature generation and verification with keys up to 4096 bits in length
- High-throughput Secure Hash Algorithm (SHA), MD5 message digest algorithm, Hash-Based Message Authentication Code (HMAC), Cipher-based Message Authentication Code (CMAC), Data Encryption Standard (DES), Triple Data Encryption Standard (Triple DES), and Advanced Encryption Standard (AES)-based encryption for data integrity assurance and confidentiality, including AES Key Wrap (AESKW) that conforms to ANSI X9.102
- Elliptic-curve cryptography (ECC) for digital signature and key agreement
- Support for smart card applications and personal identification number (PIN) processing
- Secure time-of-day
- Visa Data Secure Platform (DSP) point-to-point encryption (P2PE) with standard Visa format-preserving encryption (FPE) and format- preserving, Feistel-based Format Preserving Encryption (FF1, FF2, FF2.1). Format Preserving Counter Mode (FPCM) as defined in x9.24 Part 2
- Attributes provided: Cryptographic Coprocessor and Accelerator Functions
- Attributes required: One full-high PCIe Gen3 slot which uses a blind swap cassette
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:
  - AIX supported
  - IBM i supported
  - Linux - supported

  Linux software support can be downloaded from the following location:
  https://www.ibm.com/security/cryptocards/pciecc4/overview

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJBC) 4-NVMe U.2 (7mm) Flash drive bays

- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJRL) Non-paired Indicator EJ14 PCIe SAS RAID+ Adapter

Feature EJRL must be added for every instance of a non-paired SAS RAID adapter #EJ14. It identifies a specific high availability configuration supported by AIX or Linux which has one #EJ14 on one system and the paired #EJ14 located on a second system. IBM i does not support paired adapter on different servers.

SAS RAID adapter firmware disables write cache until a second SAS RAID adapter is recognized as its pair. This specify indicates the pairing will not be on just one server.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: Only one #EJ14 on a server and its pair on a different server.
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required: Refer to #EJ14
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No

(#EJRU) Non-paired Indicator EJ0L PCIe SAS RAID Adapter

Feature EJRU must be added for every instance of a non-paired SAS RAID adapter #EJ0L. It identifies a specific high availability configuration supported by AIX or Linux which has one #EJ0L on one system and the paired #EJ0L located on a second system. IBM i does not support paired adapter on different servers.

- Attributes provided: SAS RAID adapter firmware disables write cache until a second SAS RAID adapter is recognized.
- Attributes required: Every #EJ0L requires a 6Gb/s SAS RAID adapter (#EJ0L) on both this server and on another server that will pair up the SAS RAID adapter and enable the onboard caches to function.
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No

(#EJW1) Specify Mode-1 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)YO12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 1 (one group of 24 drive bays). It will be connected to one #EJ0J/#EL59/#EJ0M/ #EL36 PCIe3 RAID SAS adapter.

One YO12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECDT, #ECDU, #ECDV or #ECDW.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJW2) Specify Mode-1 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 1 (one group of 24 drive bays). It will be connected to two (one pair) #EJ0J/ #EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

Two YO12 cables connect the SAS Storage Enclsoure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECDT, #ECDU, #ECDV or #ECDW.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJW3) Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to two (one pair) #EJ0J/ #EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

Two X12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description

  Minimum required: 0

-
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJW4) Specify Mode-2 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to four (two pair) #EJ0J/ #EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

Two X12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJW5) Specify Mode-4 & (4)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 4 (four groups of 6 drive bays). It will be connected to four (unpaired) #EJ0J/#EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

Two X12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0

  Maximum allowed: 168 (Initial order maximum: 168)

-
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJW6) Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to one #EJ0J/#EL59/#EJ0M/ #EL36 PCIe3 RAID SAS adapter.

Two YO12 cables connect the SAS Storage Enclsoure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECDT, #ECDU, #ECDV or #ECDW.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJW7) Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (2)YO12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to two (nonpaired) #EJ0J/ #EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

Two YO12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBT, #ECBU, #ECBV or #ECBW.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  IBM i - not supported

-
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWA) Specify Mode-2 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)YO12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to one #EJ0J/#EL59/#EJ0M/ #EL36 PCIe3 RAID SAS adapter.

One YO12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBT, #ECBU, #ECBV or #ECBW.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a "partial" configuration where there are not enough adapters/cables to run all the SAS bays in the SAS Storage Enclosure. If adapters/cables are MES added later to support the rest of the SAS bays, then this specify code should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJW7.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWB) Specify Mode-2 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to two (one pair) #EJ0J/ #EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

One X12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or#ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a "partioal" configuration where there are not enough adapters/cables to run all the SAS bays in the SAS Storage Enclosure. If adapters/cables are MES added later to support the rest of the SAS bays, then this specify code should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJW4.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)

- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWC) Specify Mode-4 & (1)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 4 (four groups of 6 drive bays). It will be connected to one #EJ0J/#EL59/ #EJ0M/#EL36 PCIe3 RAID SAS adapter.

One X12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL. One leg of the X12 cable is left unattached at the adapter end.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This indicates a "partial" configuration where there are not enough adapters/cables to run all the SAS bays in the SAS Storage Enclosure. If adapters/cables are MES added later to support the rest of the SAS bays, then this specify code should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJW5, #EJWD or #EJWE. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing that in this configuration one EXP24SX Drawer (#ESLS/ELLS) should be configured in Mode 4 and be using one PCIe3 RAID SAS adapter (#EJ0J/EL59/EJ0M/ EL3B) and one X12 SAS Cables. One leg of the X cable is left unattached at the adapter end.

This indicates a "partial" configuration where there are not enough adapters/cables to run all the SAS bays in the EXP24SX. If adapters/ cables are MES added later to support the rest of the SAS bays, then this specify code should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJW5, #EJWD or #EJWE.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWD) Specify Mode-4 & (2)EJ0J/EJ0M/EL3B/EL59 & (1)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 4 (four groups of 6 drive bays). It will be connected to Two (nonpaired) #EJ0J/#EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

One X12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This indicates a "partial" configuration where there are not enough adapters/cables to run all the SAS bays in the SAS Storage Enclosure. If adapters/cables are MES added later to support the rest of the SAS bays, then this specify code

should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJW5 or #EJWE.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWE) Specify Mode-4 & (3)EJ0J/EJ0M/EL3B/EL59 & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 4 (four groups of 6 drive bays). It will be connected to Three (nonpaired) #EJ0J/#EL59/#EJ0M/#EL36 PCIe3 RAID SAS adapters.

Two X12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL. One leg of one of the two X12 cables is left unattached at the adapter end.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This indicates a "partial" configuration where there are not enough adapters/cables to run all the SAS bays in the SAS Storage Enclosure. If adapters/cables are MES added later to support the rest of the SAS bays, then this specify code should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJW5.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWF) Specify Mode-1 & (2)EJ14 & (2)YO12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 1 (one group of 24 drive bays). It will be connected to two #EJ14 PCIe3 RAID SAS adapters.

Two YO12 cables connect the SAS Storage Enclsoure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECDT, #ECDU, #ECDV or #ECDW.

One specify feature should be ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWG) Specify Mode-2 & (2)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to two #EJ14 PCIe3 RAID SAS adapters.

Two X12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWH) Specify Mode-2 & (2)EJ14 & (1)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to two #EJ14 PCIe3 RAID SAS adapters.

One X12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or#ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This indicates a "partial" configuration where there are not enough adapters/cables to run all the SAS bays in the SAS Storage Enclosure. If adapters/cables are MES added later to support the rest of the SAS bays, then this specify code should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJWG.

Attributes provided: Communicate configuration information to IBM Manufacturing

-
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWJ) Specify Mode-2 & (4)EJ14 & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to four #EJ14 PCIe3 RAID SAS adapters.

Two X12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EJWP) Specify Mode-1 & (2)EJ0L & (2)YO12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 1 (one group of 24 drive bays). It will be connected to two #EJ0L PCIe3 RAID SAS adapters.

Two YO12 cables connect the SAS Storage Enclsoure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECDT, #ECDU, #ECDV or #ECDW.

One specify feature should ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description

- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 0)
- OS level required:
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EJWR) Specify Mode-2 & (4)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to four #EJ0L PCIe3 RAID SAS adapters.

Two X12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should be ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 0)
- OS level required:
  - IBM i - not supported
  - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EJWS) Specify Mode-2 & (2)EJ0L & (2)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to two #EJ0L PCIe3 RAID SAS adapters.

Two X12 cables connect the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should be ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This specify feature indicates a full set of adapters plus SAS cables is used for the enclosure. A different specify feature code is used to indicate a "partial" configuration when a subset of adapters and cables are used.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 0)
- OS level required:
  - IBM i - not supported

    - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EJWT) Specify Mode-2 & (2)EJ0L& (1)X12 for EXP24SX #ESLS/ELLS

Feature indicates that EXP24SX SAS Storage Enclosure (#ESLS or #ELLS) will be configured by IBM Manufacturing in Mode 2 (two groups of 12 drive bays). It will be connected to two #EJ0L PCIe3 RAID SAS adapters.

One X12 cable connects the SAS Storage Enclosure to the SAS adapter ports. Use the cable length appropriate to the configuration and select from features such as #ECBJ, #ECBK or #ECBL.

One specify feature should be ordered with each SAS Storage Enclosure. This feature does not order or ship any hardware, but indicates to IBM config tools and to IBM Manufacturing the combination of enclosure mode, SAS adapter and SAS cable type which will be used.

This indicates a "partial" configuration where there are not enough adapters/cables to run all the SAS bays in the SAS Storage Enclosure. If adapters/cables are MES added later to support the rest of the SAS bays, then this specify code should be removed and the appropriate specify feature added to help IBM config tools understand the expanded usage, probably specify feature #EJWG.

- Attributes provided: Communicate configuration information to IBM Manufacturing
- Attributes required: SAS Storage Enclosure, PCIe3 RAID SAS adapter(s) and SAS cable(s) as indicated in description
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 0)
- OS level required:
    - IBM i - not supported
    - Refer to #ESLS or #ELLS to find the supported O/S levels.
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ELC0) PDU Access Cord 0.38m

This 0.38 meter (14 inch) cord is used with a vertically mounted PDU (Power Distribution Unit) such as a#ECJJ, #ECJN, #EPTJ, #EPTN, #7188 or #7109 when the PDU is located in a 7965-S42, 7965-94Y, or #ER05 Slim Rack. One end of this power cord connects to the PDU. The other end of this cord connects to the power cord running to the wall outlet or electrical power source. #ELC0 isnot supported on ECJL, ECJQ, EPTL or EPTQ PDUs.

One PDU Access Cord is required per vertically mounted PDU. Without a PDU Access Cord, inserting and removing the wall outlet power cord into the PDU can be very difficult in the narrow side pockets of the Slim Rack. A PDU Access Cord is not required for PDUs in wider racks such as the 7965-S42.

- Attributes provided: Power cord
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No

(#ELC1) 4.3m (14-Ft) PDU to Wall 24A 200-240V Power Cord North America

This power cord goes from the chassis to a wall-type outlet. The line cord has a 200-240V, IEC309 30A P+N+G Plug (de-rated to 24A) and Souriau inlet compatible with PDU FCs ECJJ and ECJN. The following countries/regions use the #ELC1 power cord to power the system: United States, Canada, Mexico.

- Attributes provided: Power cord PDU to Wall
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELC2) 4.3m (14-Ft) PDU to Wall 3PH/24A 415V Power Cord North America

This power cord goes from the chassis to a wall-type outlet. The line cord has a 3-phase wye-wired 415/240V, IEC309 30A 3P+N+G Plug (de-rated to 24A per phase) and Souriau inlet compatible with PDU FCs ECJJ and ECJN. The following countries/regions use the #ELC2 power cord to power the system: United States, Canada, Mexico.

- Attributes provided: Power Cord PDU to Wall
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 9999)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELC5) Power Cable - Drawer to IBM PDU (250V/10A)

This feature permits manufacturing to select the optimum PDU power jumper cord length (2.8M or 4.3M) for rack integration. This feature is mandatory for servers that use power supplies with C14 inlets that are going to be factory integrated with IBM racks (such as with 7014-T00 or T42 racks) that contains C19 PDU types.

Feature is not valid on initial order with non-factory integrated feature 4650. Power jumper cord has C13 on one end (for C14 power supply connector on system unit or I/O drawer) and C20 on the other end (for IBM PDU C19 receptacle). MES orders of FC #ELC5 will ship 4.3m length. If MES customers want 2.8m length should order #6665.

- Attributes provided: Power jumper cord (2.8m or 4.3m)
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EHC9) Solution Edition for Healthcare 3.7 GHZ, 60-core Processor (CEC)

- Minimum required: 0
- Maximum allowed: 2 (Initial order maximum: 2)
  OS level required:
  - IBM i 7.5, or later

    - IBM i - 7.3 TR11 or later; IBM i 7.4 TR5 or later Supported
    - AIX - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EL9D) - ESMD Load Source Specify (931GB SSD SFF-2)

This specify code indicates that a #ESMD Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESMD
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESMD
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EL9H) - ESMH Load Source Specify (1.86TB SSD SFF-2)

This specify code indicates that a #ESMH Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESMH
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESMH
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EL9S) - ESMS Load Source Specify (3.72TB SSD SFF-2)

This specify code indicates that a #ESMS Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESMS
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESMS
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EL9X) - ESMX Load Source Specify (7.44TB SSD SFF-2)

This specify code indicates that a #ESMX Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESMX

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESMX
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELCL) Power Linux processor activation for #EDP2

This feature delivers one linux processor activation for #EDP2

- Attributes provided: One linux only processor activation
- Attributes required: Linux OS on server
- Minimum required: 0
- Maximum allowed: 160 (Initial order maximum: 160)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELCM) Power Linux processor activation for #EDP4

This feature delivers one linux processor activation for #EDP4

- Attributes provided: One linux only processor activation
- Attributes required: Linux OS on server
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature ELCM not orderable in China

(#ELCN) PowerVM for Linux indicator for HEX

- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required:
  - IBM i - Not Supported
  - AIX - Not Supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELCP) 1 core Processor activation for #EHC9 no cost

For every purchased processor activation #EDPD, one free activation #ELCP will be provided. A Minimum requirement of eidht #EDPD is required.

Each occurrence of this feature will permanently activate one processor core on Processor Card #EHC9

- Attributes provided: One processor core activation for #EHC9
- Attributes required: #EHC9 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 60 (Initial order maximum: 60)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELCQ) Power Linux processor activation for #EDP3

This feature delivers one linux processor activation for #EDP3

- Attributes provided: One linux only processor activation
- Attributes required: Linux OS on server
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELKM) ESKM Load Source Specify (931GB SSD SFF-2)

This specify code indicates that a #ESKM Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESKM
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESKM
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELKR) ESKR Load Source Specify (1.86TB SSD SFF-2)

This specify code indicates that a #ESKR Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESKR
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESKR
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELKV) ESKV Load Source Specify (3.72TB SSD SFF-2)

This specify code indicates that a #ESKV Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESKV
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESKV
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELKZ) ESKZ Load Source Specify (7.44TB SSD SFF-2)

This specify code indicates that a #ESKZ Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESKZ
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESKZ
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELME) 512 GB Power Linux Memory Activations for HEX

This feature delivers 512 GB Memory linux only for HEX

- Attributes provided: 512 GM Memory linux only for HEX
- Attributes required: Linux OS on server
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELT2) #ESF2 Load Source Specify (1.1TB HDD SFF-2)

This specify code indicates that a #ESF2 disk drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: #ESF2
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - See feature ESF2

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes

- Return parts MES: No

(#ELTS) #ESFS Load Source Specify (1.7TB HDD SFF-2)

This specify code indicates that a #ESF5 disk drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: #ESFS
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - See feature ESFS

  Refer to the Software requirements section for specific O/S levels supported.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELTU) #ESEU Load Source Specify (571GB HDD SFF-2)

This specify code indicates that a #ESEU disk drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: #ESEU
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - See feature ESEU

  Refer to the Software requirements section for specific O/S levels supported.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELU9) ESK9 Load Source Specify (387GB SSD SFF-2)

This specify code indicates that a #ESK9 Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESK9
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESK9
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELUD) ESKD Load Source Specify (775GB SSD SFF-2)

This specify code indicates that a #ESKD Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESKD

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESKD
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELUH) ESKH Load Source Specify (1.55TB SSD SFF-2)

This specify code indicates that a #ESKH Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESKH
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ESKH
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELUK) ESJK Load Source Specify (931GB SSD SFF-2)

This specify code indicates that a #ESJK Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESJK
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ELUL) #ESNL Load Source Specify (283GB HDD SFF-2)

This specify code indicates that a #ESNL Hard Disk Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature #ESNL
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature ESNL
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELUM) ESJM Load Source Specify (1.86TB SSD SFF-2)

This specify code indicates that a #ESJM Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing

- Attributes required: Feature ESJM
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - See feature ESJM
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ELUP) ESJP Load Source Specify (3.72TB SSD SFF-2)

This specify code indicates that a #ESJP Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESJP
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - See feature ESJP
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ELUQ) #ESNQ Load Source Specify (571GB HDD SFF-2)

This specify code indicates that a #ESNQ Hard Disk Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature #ESNQ
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature ESNQ
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELUR) ESJR Load Source Specify (7.44TB SSD SFF-2)

This specify code indicates that a #ESJR Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESJR
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - See feature ESJR
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ELV9) - ETK9 Load Source Specify (387 GB SSD SFF-2)

This specify code indicates that a #ETK9 Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ETK9
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ETK9
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELVD) - ETKD Load Source Specify (775 GB SSD SFF-2)

This specify code indicates that a #ETKD Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ETKD
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ETKD
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELVH) - ETKH Load Source Specify (1.55 TB SSD SFF-2)

This specify code indicates that a #ETKH Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ETKH
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: See feature ETKH
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ELZ5) ES95 Load Source Specify (387GB SSD SFF-2)

This specify code indicates that a #ES95 Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ES95
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 0)
- OS level required: see feature #ES95
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ELZB) ESNB Load Source Specify (775GB SSD SFF-2)

This specify code indicates that a #ESNB Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESNB
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 0)
- OS level required: see feature #ESNB
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ELZF) ESNF Load Source Specify (1.55TB SSD SFF-2)

This specify code indicates that a #ESNF Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESNF
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 0)
- OS level required: see feature #ESNF
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EM8F) Active Memory expansion enablement for HEX

Feature number EM8F instructs manufacturing to enable memory expansion on the system.

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - IBM i - not supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EMAZ) 1 GB Memory activation for HEX

1 GB Memory Activation

- Attributes provided: 1 GB Activation Code Number
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 65536 (Initial order maximum: 65536)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EMBK) 500 GB DDR4 Mobile Memory Activation for HEX

One #EMBK feature is ordered for each 500 GB mobile memory activation for HEX

- Attributes provided: Enabled activation of 500 GB mobile memory for HEX.
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 130 (Initial order maximum: 130)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EMBZ) 512 GB Memory Activations for HEX

Each occurrence of this feature will permanently activate 512 GB of DDR4, Power10 memory

- Attributes provided: Quantity 512 of 1GB Memory Activation.
- Attributes required: Inactive DDR4 Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 128 (Initial order maximum: 128)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EMC1) 128 GB (4x32GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory

Feature ships four 32 GB DDIMMs providing 128 GB of DDR4 memory. DRAM is 16Gb and 3200 MHz.

Limitations:

- Only one size DDIMM can be placed on one SCM. (multiple size DDIMM can be placed on one node)
- Attributes provided: Four 32 GB DDIMMs providing 128 GB of DDR4 memory
- Attributes required: Four empty DDIMM slots
- Minimum required: 0
- Maximum allowed: 64 (Initial order maximum: 64)
- OS level required:
  - IBM i - Supported
  - AIX - Supported
  - Linux - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EMC2) 256 GB (4x64GB) DDIMMs, 3200 MHz, 16GBIT DDR4 Memory

Feature ships four 64 GB DDIMMs providing 256 GB of DDR4 memory. DRAM is 16Gb and 3200 MHz.

Limitations:

- Only one size DDIMM can be placed on one SCM. (multiple size DDIMM can be placed on one node)
- Attributes provided: Four 64 GB DDIMMs providing 256 GB of DDR4 memory
- Attributes required: Four empty DDIMM slots

- Minimum required: 0
- Maximum allowed: 64 (Initial order maximum: 64)
- OS level required:
  - IBM i - Supported
  - AIX - Supported
  - Linux - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No

(#EMC3) 512 GB (4x128GB) DDIMMs, 2933 MHz, 16GBIT DDR4 Memory

Feature ships four 128 GB DDIMMs providing 512 GB of DDR4 memory. DRAM is 16Gb and 2933 MHz.

Limitations:

- Only one size DDIMM can be placed on one SCM. (multiple size DDIMM can be placed on one node)
- Attributes provided: Four 128 GB DDIMMs providing 512 GB of DDR4 memory.
- Attributes required: Four empty DDIMM slots
- Minimum required: 0
- Maximum allowed: 64 (Initial order maximum: 64)
- OS level required:
  - IBM i - Supported
  - AIX - Supported
  - Linux - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No
  Note: Genaral availability date December 10, 2021.

(#EMC4) 1 TB (4x256GB) DDIMMs, 2933 MHz, 16GBIT DDR4 Memory

Feature ships four 256 GB DDIMMs providing 1 TB of DDR4 memory. DRAM is 16Gb and 2933 MHz.

Limitations:

- Only one size DDIMM can be placed on one SCM. (multiple size DDIMM can be placed on one node)
- Attributes provided: Four 256 GB DDIMMs providing 1 TB of DDR4 memory
- Attributes required: Four empty DDIMM slots
- Minimum required: 0
- Maximum allowed: 64 (Initial order maximum: 64)
- OS level required:
  - IBM i - Supported
  - AIX - Supported
  - Linux - Supported
- Initial Order/MES/Both/Supported: Both
- CSU: No
- Return parts MES: No
  Note: Genaral availability date December 10, 2021.

(#EMAC) 512 GM Memory Activation for #EHC9

For every purchased memory activation #EMBZ, one free activation #EMAC will be provided. A Minimum of one #EMBZ is required. Additional memory activations may be purchased using normal priced static activations.

This feature delivers 512 GB Memory activation for #EHC9

- Attributes provided: 512 GB Memory Activation for #EHC9
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 8 (Initial order maximum: 8)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EME0) 100GB Static to Mobile Memory Auto Conversion

Feature is used to order automatic conversion of a static action to a mobile activation without IBM intervention. The term "static activations" also includes Mobile Enabled activations which are technically static. Firmware 840 or later is a pre-requisite. Use of this feature on servers with firmware 830 or earlier will require IBM intervention through the pcod@us.ibm.com

- Attributes provided: Automated conversion
- Attributes required: Firmware 840
- Minimum required: 0
- Maximum allowed: 246 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No

(#EMQZ) 100 GB of #EMAZ Memory activation for HEX

This feature delivers 100 GB of Feature #EMAZ Memory activation at no cost

- Attributes provided: 100 GB Memory Activation
- Attributes required: Inactive Memory available for temporary use installed
- Minimum required: 0
- Maximum allowed: 655 (Initial order maximum: 655)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EMX0) PCIe Gen3 I/O Expansion Drawer

This 19-inch, 4U (4 EIA) enclosure provides PCIe Gen3 slots outside of the system unit. It has two module bays. One 6-Slot Fanout Module (#EMXH) can be placed in each module bay. Two 6-slot modules provide a total of 12 PCIe Gen3 slots. Each fanout module is connected to a PCIe3 Optical Cable Adapter located in the system unit over an active optical CXP cable (AOC) pair or CXP copper cable pair.

The PCIe Gen3 I/O Expansion Drawer has two redundant, hotplug power supplies. Each power supply has its own separately ordered power cord. The two power cords plug into a Power supply conduit which connects to the power supply. The single-phase AC power supply is rated at 1030 Watt and can use 100-120V or 200-240V. If using 100-120V, then the maximum is 950 Watt. It's recommended the power supply connect to a PDU in the rack. Power Systems PDUs are designed for 200-240V electrical source.

The drawer has fixed rails which can accommodate racks with depths from 27.5 inches to 30.5 inches.

Limitations:

- #EMX0 has a cable management bracket located at the rear of the drawer which swings up to provide service access to the PCIe adapters. 2U (2 EIA) of space is required to swing up the bracket. Thus the drawer cannot be placed in the very top 2U of a rack.
- There is a power cord access consideration with vertically mounted PDUs on the right hand side of the rack when viewed from the rear of the rack. The #EMX0 cable management bracket makes accessing some of the PDU outlets located at the same rack height as the #EMX0 drawer more challenging. Using a horizontally mounted PDU or locating the PDU or #EMX0 at a different vertical location is recommended.
- Attributes provided: 19-inch 4U (4 EIA) PCIe Gen3 I/O Expansion Drawer
- Attributes required: One or two PCIe Optical Cable Adapters (#EJ24), one or two PCIe3 fanout modules (#EMXH), one or two CXP cable pairs (such as #ECCR or #ECCY), one power supply conduit (such as #EMXA).
- Minimum required: 0
- Maximum allowed: 16 (Initial order maximum: 16)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to the Software requirements section for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EMXA) AC Power Supply Conduit for PCIe3 Expansion Drawer

Provides two 320-C14 inlet electrical connections for two separately ordered AC power cords with C13 connector plugs. Conduit provides electrical power connection between two power supplies located in the front of a PCIe Gen3 I/O Expansion Drawer (#EMX0) and two power cords which connect on the rear of the PCIe Gen3 I/O Expansion Drawer.

- Attributes provided: Two AC Power Supply connections
- Attributes required: PCIe Gen3 I/O Expansion Drawer and two AC power cords
- Minimum required: 0
- Maximum allowed: 16 (Initial order maximum: 16)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EMXH) PCIe3 6-Slot Fanout Module for PCIe3 Expansion Drawer

PCIe3 fanout module for PCIe Gen3 I/O Expansion Drawer. Provides six PCIe Gen3 full high, full length slots (two x16 and four x8). The PCIe slots are hot plug.

The module has two CXP ports which are connected two CXP ports on a PCIe Optical Cable Adapter (only allowed to connect to #EJ24 Optical Cable Adapter). A pair of active optical CXP cables (AOC) or a pair of CXP copper cables are used for this connection. The top CXP port of the fanout module is cabled to the top CXP port of the PCIe3 Optical Cable Adapter. The bottom CXP port of the fanout module is cabled to the bottom CXP port of the same PCIe3 Optical Cable Adapter.

Limitations:

- Mixing of prior PCIe3 fanout modules (#EMXF, #EMXG, #ELMF, #ELMG) with PCIe3 fanout module (feature #EMXH) in the same I/O Expansion Drawer is not allowed

  PCIe3 Optical Cable Adapters (#EJ24) requires to use Optical Cables (#ECCR, ECCY, or #ECCZ)

-
- Attributes provided: PCIe3 6-slot fanout module for PCIe Gen3 I/O Expansion Drawer
- Attributes required: Available bay in PCIe Gen3 I/O Expansion Drawer.
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - IBM i supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  - Firmware 9.20, or later for copper CXP cables

(#EN01) 1m (3.3-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper

1m (3.3-ft) copper active twinax Ethernet cable which supports Ethernet data transfer rates up to 10 Gb/s. The cable has a copper twinax transceiver on each end which is placed in an SFP+ port of an adapter and/or a switch. This cabling option can be a cost effective alternative to optical cable for short reach link high-speed connection.

- Attributes provided: 10Gb/s copper active twinax Ethernet cable
- Attributes required: One available SFP+ 10Gb/s Ethernet Port
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EN02) 3m (9.8-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper

3m (9.8-ft) copper active twinax Ethernet cable which supports Ethernet data transfer rates up to 10 Gb/s. The cable has a copper twinax transceiver on each end which is placed in an SFP+ port of an adapter and/or a switch. This cabling option can be a cost effective alternative to optical cable for short reach link high-speed connection.

- Attributes provided: 10Gb/s copper active twinax Ethernet cable
- Attributes required: One available SFP+ 10Gb/s Ethernet Port
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EN03) 5m (16.4-ft), 10Gb E'Net Cable SFP+ Act Twinax Copper

5m (16.4-ft) copper active twinax Ethernet cable which supports Ethernet data transfer rates up to 10 Gb/s. The cable has a copper twinax transceiver on each end which is placed in an SFP+ port of an adapter and/or a switch. This cabling option can be a cost effective alternative to optical cable for short reach link high-speed connection.

- Attributes provided: 10Gb/s copper active twinax Ethernet cable
- Attributes required: One available SFP+ 10Gb/s Ethernet Port

- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EN0S) PCIe2 4-Port (10Gb+1GbE) SR+RJ45 Adapter

PCIe Gen2 x8 short Ethernet adapter supports Ethernet NIC (Network Interface Card) traffic. The adapter provides two 10 Gb SR optical ports and two 1 Gb RJ45 ports in a PCIe 8x short form adapter.

The 10Gb ports are SFP+ and include optical SR transceivers. The ports have LC Duplex type connectors and utilize shortwave laser optics and MMF-850nm fiber cabling. With 62.5 micron OM1, up to 33 meter length fiber cables are supported. With 50 micron OM2, up to 82 meter fiber cable lengths are supported. With 50 micron OM3 or OM4, up to 300 meter fiber cable lengths are supported.

For the 1Gb RJ45 ports, 4-pair CAT-5 Unshielded Twisted Pair (UTP) cable or higher is supported for distances of up to 100 meters. In addition to 1Gb (1000 Mb) networks, 100 Mb networks are also supported, but 10Mb networks are not supported.

#EN0S and #EN0T adapters are electronically identical. They are physically identical except #EN0S has a tail stock for full high PCIe slots and #EN0T has a tail stock for low profile slots. The CCIN is 2CC3 for both features.

Details for the ports include:

- VIOS NIM and LINUX NETWORK INSTALL are supported.
- IEEE 802.3ae (10GBASE-SR), IEEE 802.3ab (1000BASE-T GbE), IEEEu 802.3u (100BASE-T), 802.1p priority, 802.1Q VLAN tagging, 802.3x flow control, 802.3ad load-balancing and failover,
- Link aggregation, IEEE 802.3ad 802.3
- Multiple MAC addresses/promiscuous mode (for PowerVM/VIOS) per interface
- Message Signalling Interrupt MSI-X, MSI and support of legacy pin interrupts
- Ether II and IEEE 802.3 encapsulated frames
- Jumbo frames up to 9.6 Kbytes
- TCP checksum offload for IPv4 and IPv6
- TCP segmentation Offload (TSO) RSS (Receive Side Scaling) support for IPv4, IPv6 and UDP for IPv4 and IPv6
- UDP checksum offload for IPv4 and IPv6
- AIX, IBM i, and Linux provide software iSCSI support through the adapter.
- Attributes provided: Four ports - two 10Gb and two 1Gb E
- Attributes required: PCIe Gen2 or Gen3 slot - full high
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 0)
- OS level required:
  - AIX - supported
  - IBM i supported through VIOS
  - Linux - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No
  - VIOS - supported
  - AIX NIM and Linux Network Install are supported.
  - AIX, IBM i and Linux provide software iSCSI support through the adapter.

(#EN0T) PCIe2 LP 4-Port (10Gb+1GbE) SR+RJ45 Adapter

PCIe Gen2 x8 short Ethernet adapter supports Ethernet NIC (Network Interface Card) traffic. The adapter provides two 10 Gb SR optical ports and two 1 Gb RJ45 ports in a PCIe 8x short form adapter.

The 10Gb ports are SFP+ and include optical SR transceivers. The ports have LC Duplex type connectors and utilize shortwave laser optics and MMF-850nm fiber cabling. With 62.5 micron OM1, up to 33 meter length fiber cables are supported. With 50 micron OM2, up to 82 meter fiber cable lengths are supported. With 50 micron OM3 or OM4, up to 300 meter fiber cable lengths are supported.

For the 1Gb RJ45 ports, 4-pair CAT-5 Unshielded Twisted Pair (UTP) cable or higher is supported for distances of up to 100 meters. In addition to 1Gb (1000 Mb) networks, 100 Mb networks are also supported, but 10Mb networks are not supported.

#EN0S and #EN0T adapters are electronically identical. They are physically identical except #EN0S has a tail stock for full high PCIe slots and #EN0T has a tail stock for low profile slots. The CCIN is 2CC3 for both features.

Details for the ports include:

- VIOS NIM and LINUX NETWORK INSTALL are supported.
- IEEE 802.3ae (10GBASE-SR), IEEE 802.3ab (1000BASE-T GbE), IEEEu 802.3u (100BASE-T), 802.1p priority, 802.1Q VLAN tagging, 802.3x flow control, 802.3ad load-balancing and failover,
- Link aggregation, IEEE 802.3ad 802.3
- Message Signalling Interrupt MSI-X, MSI and support of legacy pin interrupts
- Ether II and IEEE 802.3 encapsulated frames
- Jumbo frames up to 9.6 Kbytes
- TCP checksum offload for IPv4 and IPv6
- TCP segmentation Offload (TSO) RSS (Receive Side Scaling) support for IPv4, IPv6 and UDP for IPv4 and IPv6
- UDP checksum offload for IPv4 and IPv6
- AIX, IBM i and Linux provide software iSCSI support through the adapter.
- Attributes provided: Four ports - two 10Gb and two 1Gb E
- Attributes required: PCIe Gen2 or Gen3 slot - low profile
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 0)
- OS level required:
  - AIX supported
  - IBM i is supported only with VIOS
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No
  Note: VIOS supported

(#EN0X) PCIe2 LP 2-port 10/1GbE BaseT RJ45 Adapter

PCIe Gen2 short x8 adapter which provides two 10G-BaseT ports. The ports are RJ45. The ports default to auto negotiate the highest speed either 10Gb (10GBaseT), 1Gb (1000BaseT) or 100Mb (100BaseT) full duplex. Each RJ45 port's configuration is independent of the other. The adapter supports Ethernet NIC (Network Interface Card) traffic.

The RJ45 ports use 4-pair CAT-6A cabling for distances of up to 100 meters or CAT-6 cabling for distances up to 37 meters. CAT5 cabling is not tested and is not supported.

#EN0W and #EN0X are electronically identical with the same CCIN of 2CC4. #EN0W has a full high tail stock and #EN0X has a low profile tail stock.

Details for the ports include:

- NIM install supported for VIOS, AIX, and Linux
- IEEE 802.3an (10GBASE-T), IEEE 802.3ab (1000BASE-T GbE), IEEEu (100BASE-T), 802.1p priority, 802.1Q VLAN tagging, 802.3x flow control, 802.3ad load-balancing and failover, Link aggregation, IEEE 802.3ad 802.3
- Multiple MAC addresses / promiscuous mode (for PowerVM/VIOS) per interface
- Message Signalling Interrupt (MSI-X, MSI) support of legacy pin interrupts
- Ether II and IEEE 802.3 encapsulated frames
- Jumbo frames up to 9.6 Kbytes
- TCP checksum offload for IPv4 and IPv6
- TCP segmentation Offload (TSO)
- RSS (Receive Side Scaling) support for IPv4, IPv6 and UDP
- UDP checksum offload for IPv4 and IPv6
- AIX, IBM i, and Linux provide software iSCSI support through the adapter
- Attributes provided: Two 10G-BaseT ports
- Attributes required: PCIe Gen2 or Gen3 slot - low profile
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - AIX - supported
  - IBM i supported only with VIOS
  - Linux - supported
  - Refer to the Software Requirements section to find the supported O/S levels for AIX and Linux.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: VIOS supported

(#EN0W) PCIe2 2-port 10/1GbE BaseT RJ45 Adapter

PCIe Gen2 short x8 adapter which provides two 10G-BaseT ports. The ports are RJ45. The ports default to auto negotiate the highest speed either 10Gb (10GBaseT), 1Gb (1000BaseT) or 100Mb (100BaseT) full duplex. Each RJ45 port's configuration is independent of the other. The adapter supports Ethernet NIC (Network Interface Card) traffic.

The RJ45 ports use 4-pair CAT-6A cabling for distances of up to 100 meters or CAT-6 cabling for distances up to 37 meters. CAT5 cabling is not tested and is not supported.

#EN0W and #EN0X are electronically identical with the same CCIN of 2CC4. #EN0W has a full high tail stock and #EN0X has a low profile tail stock.

Details for the ports include:

- NIM install supported for VIOS, AIX, and Linux
- IEEE 802.3an (10GBASE-T), IEEE 802.3ab (1000BASE-T GbE), IEEEu (100BASE-T), 802.1p priority, 802.1Q VLAN tagging, 802.3x flow control, 802.3ad load-balancing and failover, Link aggregation, IEEE 802.3ad 802.3
- Multiple MAC addresses / promiscuous mode (for PowerVM/VIOS) per interface
- Message Signalling Interrupt (MSI-X, MSI) support of legacy pin interrupts
- Ether II and IEEE 802.3 encapsulated frames
- Jumbo frames up to 9.6 Kbytes
- TCP checksum offload for IPv4 and IPv6
- TCP segmentation Offload (TSO)
- RSS (Receive Side Scaling) support for IPv4, IPv6 and UDP
- UDP checksum offload for IPv4 and IPv6

  AIX, IBM i, and Linux provide software iSCSI support through the adapter

-
- Attributes provided: Two 10G-BaseT ports
- Attributes required: PCIe Gen2 or Gen3 slot - full high
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:
  - AIX - supported
  - IBM i supported only with VIOS
  - Linux - supported
  - Refer to the Software Requirements section to find the supported O/S levels for AIX and Linux.
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: VIOS supported

(#EN1A) PCIe3 32Gb 2-port Fibre Channel Adapter

PCIe Gen3 32 Gigabit dual-port Optical Fibre Channel (FC) Adapter is a high-performance 8x short form adapter based on the Broadcom LPe32000-series PCIe Host Bus Adapter (HBA). The adapter provides two ports of 32Gb Fibre Channel capability using SR optics. Each port can provide up to 32Gb Fibre Channel functions simultaneously.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The ports are SFP+ and include an optical SR transceiver. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 8, 16, and 32Gbps and will automatically negotiate to the highest speed possible.

Feature #EN1A and #EN1B are electronically identical. They differ physically only that EN1A has a tail stock for full high PCIe slots and #EN1B has a short tail stock for low profile PCIe slots.

CCIN is 578F for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables:

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 2-port 32Gb Optical FC
- Attributes required: Full high profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:
  - IBM i - supported
  - AIX - supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EN1B) PCIe3 LP 32Gb 2-port Fibre Channel Adapter

PCIe Gen3 32 Gigabit dual-port Optical Fibre Channel (FC) Adapter is a high-performance 8x short form adapter based on the Broadcom LPe32000-series PCIe Host Bus Adapter (HBA). The adapter provides two ports of 32Gb Fibre Channel capability using SR optics. Each port can provide up to 32Gb Fibre Channel functions simultaneously.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The ports are SFP+ and include an optical SR transceiver. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 8, 16 and 32Gbps and will automatically negotiate to the highest speed possible.

Feature #EN1A and #EN1B are electronically identical. They differ physically only that EN1A has a tail stock for full high PCIe slots and #EN1B has a short tail stock for low profile PCIe slots.

CCIN is 578F for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables:

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 2-port 32Gb Optical FC

  Attributes required: Low profile PCIe Gen3 slot

-
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - IBM i - supported
  - AIX - supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported

(#EN1C) PCIe3 16Gb 4-port Fibre Channel Adapter

PCIe Gen3 16 Gigabit quad-port Optical Fibre Channel (FC) Adapter is a high-performance x8 short form PCIe adapter based on the Emulex LPe31004 PCIe Host Bus Adapter (HBA). The adapter provides four ports of 16Gb Fibre Channel capability using SR optics. Each port can provide up to 16Gb Fibre Channel functions simultaneously.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The ports are SFP+ and include an optical SR transceiver. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 4, 8 and 16 Gbps and will automatically negotiate to the highest speed possible.

Feature #EN1C and #EN1D are electronically identical. They differ physically only that EN1C has a tail stock for full high PCIe slots and #EN1D has a tail stock for low profile PCIe slots.

CCIN is 578E for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 4-port 16Gb Optical FC
- Attributes required: Full high profile PCIe Gen3 slot
- Minimum required: 0

- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:
  - IBM i - supported
  - AIX - supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported

(#EN1D) PCIe3 LP 16Gb 4-port Fibre Channel Adapter

PCIe Gen3 16 Gigabit quad-port Optical Fibre Channel (FC) Adapter is a high-performance x8 short form PCIe adapter based on the Emulex LPe31004 PCIe Host Bus Adapter (HBA). The adapter provides four ports of 16Gb Fibre Channel capability using SR optics. Each port can provide up to 16Gb Fibre Channel functions simultaneously.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The ports are SFP+ and include an optical SR transceiver. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 4, 8 and 16 Gbps and will automatically negotiate to the highest speed possible.

Feature #EN1C and #EN1D are electronically identical. They differ physically only that EN1C has a tail stock for full high PCIe slots and #EN1D has a tail stock for low profile PCIe slots.

CCIN is 578E for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 4-port 16Gb Optical FC
- Attributes required: Low profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  IBM i - supported

-
- AIX - supported
- Linux - supported

Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported

(#EN1E) PCIe3 16Gb 4-port Fibre Channel Adapter

This PCIe Gen3 16 Gigabit quad-port optical fibre channel (FC) adapter is a high-performance x8 short form adapter based on the Marvell QLE2694L PCIe host bus adapter (6.6 inches x 2.371 inches). The adapter provides four ports of 16Gb fibre channel capability using SR optics. Each port can provide up to 3,200MBps bandwidth per port.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The adapter ships with Soldered Small Form Factor (SFF) optical transceivers installed. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 4, 8, and 16Gbps and will automatically negotiate to the highest speed possible. The adapter supports boot on IBM power with fcode.

Feature #EN1E and #EN1F are electronically identical. They differ physically only that the #EN1E has a tail stock for full height PCIe slots and the #EN1F has a short tail stock for low profile PCIe slots.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 4-port 16Gb Optical FC
- Attributes required: Full high profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:
  - IBM i supported through VIOS
  - AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-03-2136 or later
  - Red Hat Enterprise Linux 8 for Power LE, version 8.4, or later
  - Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8 for Power LE version 8.4, 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

Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported

(#EN1F) PCIe3 LP 16Gb 4-port Fibre Channel Adapter

This PCIe Gen3 16 Gigabit quad-port optical fibre channel (FC) adapter is a high-performance x8 short form adapter based on the Marvell QLE2694L PCIe host bus adapter (6.6 inches x 2.371 inches). The adapter provides four ports of 16Gb fibre channel capability using SR optics. Each port can provide up to 3,200MBps bandwidth per port.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The adapter ships with Soldered Small Form Factor (SFF) optical transceivers installed. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 4, 8, and 16Gbps and will automatically negotiate to the highest speed possible. The adapter supports boot on IBM power with fcode.

Feature #EN1E and #EN1F are electronically identical. They differ physically only that the #EN1E has a tail stock for full height PCIe slots and the #EN1F has a short tail stock for low profile PCIe slots.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 4-port 16Gb Optical FC
- Attributes required: Low profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-03-2136 or later
  - IBM i supported through VIOS
  - Red Hat Enterprise Linux 8.4, for Power LE, or later
  - Red Hat Enterprise Linux for SAP, with Red Hat Enterprise Linux 8.4, 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
  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Assignment to the VIOS supported

(#EN1G) PCIe3 2-Port 16Gb Fibre Channel Adapter

The PCIe3 x8 dual-port Fibre Channel (16 Gb/s) adapter is a PCI Express (PCIe) generation 3 (Gen3) x8 adapter. This PCIe adapter is based on the Marvell QLE2692 PCIe host bus adapter 15.2 cm x 7 cm (6.6 inches x 2.7 inches). The adapter provides two ports of 16 Gb Fibre Channel capability by using SR optics. Each port can provide up to 3,200 Mbps bandwidth per port. Each port provides single initiator capability over a fiber optic link or with N_Port ID Virtualization (NPIV) it provides multiple initiator capabilities. The ports are SFP+ and include an optical SR transceiver. The ports have LC type connectors and use shortwave laser optics. The adapter operates at link speeds of 4, 8, and 16 Gbps and it automatically negotiates to the highest speed possible. The adapter supports start up on IBM Power Systems with FCode.

Feature #EN1G and #EN1H are electronically identical. They differ physically only that the #EN1G has a tail stock for full height PCIe slots and the #EN1H has a low profile tail stock.

CCIN is 579B for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables:

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: N/A o 32GFC: N/A

Note: The hardware cannot detect the type and length of cable that is installed. The link auto-negotiates to the speed that is reported during negotiation by the target. You must manually set the maximum negotiation speed. If the speed value is set higher than the supported cable value, bit errors can occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

Limitations: This adapter is not supported in India.

- Attributes provided: Enhanced diagnostics and manageability, Unparalleled performance and more efficient port utilization, Single initiator capability over a fiber optic link or with NPIV, 16 Gb/s of throughput per port, Multiple initiator capabilities.
- Attributes required: Full high profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:

- IBM i not supported and IBM i not supported with VIOS
- AIX - supported
- Red Hat Enterprise Linux 8 for Power LE, version 8.4, or later
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8 for Power LE version 8.4, 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

Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EN1H) PCIe3 LP 2-Port 16Gb Fibre Channel Adapter

The PCIe3 x8 dual-port Fibre Channel (16 Gb/s) adapter is a PCI Express (PCIe) generation 3 (Gen3) x8 adapter. This PCIe adapter is based on the Marvell QLE2692 PCIe host bus adapter 15.2 cm x 7 cm (6.6 inches x 2.7 inches). The adapter provides two ports of 16 Gb Fibre Channel capability by using SR optics. Each port can provide up to 3,200 Mbps bandwidth per port. Each port provides single initiator capability over a fiber optic link or with N_Port ID Virtualization (NPIV) it provides multiple initiator capabilities. The ports are SFP+ and include an optical SR transceiver. The ports have LC type connectors and use shortwave laser optics. The adapter operates at link speeds of 4, 8, and 16 Gbps and it automatically negotiates to the highest speed possible. The adapter supports start up on IBM Power Systems with FCode.

Feature #EN1G and #EN1H are electronically identical. They differ physically only that the #EN1G has a tail stock for full height PCIe slots and the #EN1H has a low profile tail stock.

CCIN is 579B for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

Cables:

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: N/A o 32GFC: N/A

Note: The hardware cannot detect the type and length of cable that is installed. The link auto-negotiates to the speed that is reported during negotiation by the target. You must manually set the maximum negotiation speed. If the speed value is set higher than the supported cable value, bit errors can occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

Limitations: This adapter is not supported in India.

- Attributes provided: Enhanced diagnostics and manageability, Unparalleled performance and more efficient port utilization, Single initiator capability over a fiber optic link or with NPIV, 16 Gb/s of throughput per port, Multiple initiator capabilities.

- Attributes required: Low profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - AIX - supported
  - IBM i not supported and IBM i not supported through VIOS
  - Red Hat Enterprise Linux 8.4, for Power LE, or later
  - Red Hat Enterprise Linux for SAP, with Red Hat Enterprise Linux 8.4, 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

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EN1J) PCIe4 32Gb 2-port Optical Fibre Channel Adapter

This PCIe 4.0 x8 dual-port 32 Gigabit optical fibre channel (FC) adapter is a high-performance short form adapter based on the Marvell QLE2772 PCIe host bus adapter (6.6 inches x 2.731 inches). The adapter provides two ports of 32Gb fibre channel capability using SR optics. Each port can provide up to 6,400MBps bandwidth per port.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The adapter ships with 32Gb SR optical transceivers installed. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 8, 16 and 32Gbps and will automatically negotiate to the highest speed possible. The adapter supports boot on IBM power with fcode.

Feature #EN1J and #EN1K are electronically identical. They differ physically only that the #EN1J has a tail stock for full height PCIe slots and the #EN1K has a short tail stock for low profile PCIe slots.

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 2-port 32Gb Optical FC
- Attributes required: Full high profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
  OS level required:
  - IBM i 7.5, or later

- IBM i 7.4 TR5, or later
- IBM i through VIOS - supported
- AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-03-2136 or later
- Red Hat Enterprise Linux 8 for Power LE, version 8.4, or later
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8 for Power LE version 8.4, 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

Refer to Software Requirements for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported
  Note: native IBM i support on 7.4 requires these PTFs or later:
  - MF69081
  - MF69082
  - MF69084

(#EN1K) PCIe4 LP 32Gb 2-port Optical Fibre Channel Adapter

This PCIe 4.0 x8 dual-port 32 Gigabit optical fibre channel (FC) adapter is a high-performance short form adapter based on the Marvell QLE2772 PCIe host bus adapter (6.6 inches x 2.731 inches). The adapter provides two ports of 32Gb fibre channel capability using SR optics. Each port can provide up to 6,400MBps bandwidth per port.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The adapter ships with 32Gb SR optical transceivers installed. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 8, 16 and 32Gbps and will automatically negotiate to the highest speed possible. The adapter supports boot on IBM power with fcode.

Feature #EN1J and #EN1K are electronically identical. They differ physically only that the #EN1J has a tail stock for full height PCIe slots and the #EN1K has a short tail stock for low profile PCIe slots.

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 2-port 32Gb Optical FC
- Attributes required: Low profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
  OS level required:

- IBM i 7.5, or later
- IBM i 7.4 TR5, or later
- IBM i through VIOS - supported
- AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-03-2136 or later
- Red Hat Enterprise Linux 8 for Power LE, version 8.4, or later
- Red Hat Enterprise Linux for SAP with Red Hat Enterprise Linux 8 for Power LE version 8.4, 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

Refer to Software Requirements for specific O/S levels supported.

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported
  Note: native IBM i support on 7.4 requires these PTFs or later:
  - MF69081
  - MF69082
  - MF69084

(#EN2A) PCIe3 16Gb 2-port Fibre Channel Adapter

PCIe Gen3 16 Gigabit dual-port Optical Fibre Channel (FC) Adapter is a high-performance 8x short form adapter PCIe Host Bus Adapter (HBA). The adapter provides two ports of 16Gb Fibre Channel capability using SR optics. Each port can provide up to 16Gb Fibre Channel functions simultaneously.

Each port provides single initiator capability over a fibre link. If you are using N_Port ID Virtualization (NPIV), multiple initiator capability is provided. The ports are SFP+ and include an optical SR transceiver. The ports have little connector-type (LC) and utilize shortwave laser optics. The adapter operates at link speeds of 4, 8, and 16 Gbps and it will automatically negotiate to the highest speed possible.

The adapter connects to a Fibre Channel switch at 4 Gb, 8 Gb, or 16 Gb. It can directly attach to a device without a switch at 16 Gb. The adapter without a Fibre Channel switch attached is not supported at 4 Gb or 8 Gb.

NPIV capability is supported through Virtual I/O Server (VIOS).

Feature #EN2A and #EN2B are electronically identical. They differ physically only that EN2A has a tail stock for full high PCIe slots and #EN2B has a short tail stock for low profile PCIe slots.

CCIN is 579D for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

The adapter has the following features:

- The adapter is compliant with the PCIe base and Card Electromechanical (CEM) 2.0 specifications with the following characteristics:
  - Provides an x8 lane link interface at 14.025 Gbps, 8.5 Gbps, or 4.25 Gbps (automatic negotiation with system)
  - Provides support for one Virtual Channel (VC0) and one Traffic Class (TC0)
  - Provides configuration and I/O memory read and write, completion, and messaging capabilities
  - Provides support for 64-bit addressing
  - Provides error correction code (ECC) and error protection functions
  - Provides link cyclic redundancy check (CRC) on all PCIe packets and message information
  - Provides a large payload size of 2048 bytes for read and write functions

    - Provides a large read request size of 4096 bytes
- The adapter is compatible with 4, 8, and 16 Gb Fibre Channel interface with the following characteristics:
    - Provides for automatic negotiation between 4 Gb, 8 Gb, or 16 Gb link attachments
    - Provides support for the following Fibre Channel topologies: point-to-point (16 Gb only) and fabric
    - Provides support for Fibre Channel class 3
    - Provides a maximum Fibre Channel throughput that is achieved by using full duplex hardware support
- The adapter provides an end-to-end data path parity and CRC protection, including internal data path random-access memory (RAM)
- Provides architectural support for multiple upper layer protocols
- Provides comprehensive virtualization capabilities with support for N_Port ID Virtualization (NPIV) and virtual fabric (VF)
- Provides support for message signaled interrupts extended (MSI-X)
- Provides support for 255 VFs and 1024 MSi-X
- Provides an internal, high-speed static random-access memory (SRAM) memory
- Provides ECC protection of local memory that includes single-bit correction and double-bit protection
- Provides an embedded shortwave optical connection with diagnostics capability
- Provides support for an on-board context management by firmware:
    - Up to 8192 FC port logins
    - I/O multiplexing down to the Fibre Channel frame level
- Provides data buffers capable of supporting 64+ buffer-to-buffer (BB) credits per port for shortwave applications
- Provides link management and recovery that is handled by firmware
- Provides on-board diagnostic capability accessible by an optional connection
- Provides a performance up to 16 Gbps full duplex

Cables:

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 2-port 16Gb Optical FC
- Attributes required: Full high profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required:
    - AIX supported
    - IBM i supported
    - Linux supported

Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EN2B) PCIe3 LP 16Gb 2-port Fibre Channel Adapter

PCIe Gen3 16 Gigabit dual-port Optical Fibre Channel (FC) Adapter is a high-performance 8x short form adapter PCIe Host Bus Adapter (HBA). The adapter provides two ports of 16Gb Fibre Channel capability using SR optics. Each port can provide up to 16Gb Fibre Channel functions simultaneously.

Each port provides single initiator capability over a fibre link. If you are using N_Port ID Virtualization (NPIV), multiple initiator capability is provided. The ports are SFP+ and include an optical SR transceiver. The ports have little connector-type (LC) and utilize shortwave laser optics. The adapter operates at link speeds of 4, 8, and 16 Gbps and it will automatically negotiate to the highest speed possible.

The adapter connects to a Fibre Channel switch at 4 Gb, 8 Gb, or 16 Gb. It can directly attach to a device without a switch at 16 Gb. The adapter without a Fibre Channel switch attached is not supported at 4 Gb or 8 Gb.

NPIV capability is supported through Virtual I/O Server (VIOS).

Feature #EN2A and #EN2B are electronically identical. They differ physically only that EN2A has a tail stock for full high PCIe slots and #EN2B has a short tail stock for low profile PCIe slots.

CCIN is 579D for both features.

Each port has two LED indicators located on the bracket next to each connector. These LEDs communicate boot status and give a visual indication of the operating state. The LEDs have five defined states; solid on, solid off, slow blink, fast blink, and flashing. The slow blink rate is 1Hz, the fast blink is 4Hz, and the flashing refers to an irregular on/off transition that reflects test progress. The operator should observe the LED sequence for several seconds to ensure that the operating state is correctly identified.

The adapter has the following features:

- The adapter is compliant with the PCIe base and Card Electromechanical (CEM) 2.0 specifications with the following characteristics:
  - Provides an x8 lane link interface at 14.025 Gbps, 8.5 Gbps, or 4.25 Gbps (automatic negotiation with system)
  - Provides support for one Virtual Channel (VC0) and one Traffic Class (TC0)
  - Provides configuration and I/O memory read and write, completion, and messaging capabilities
  - Provides support for 64-bit addressing
  - Provides error correction code (ECC) and error protection functions
  - Provides link cyclic redundancy check (CRC) on all PCIe packets and message information
  - Provides a large payload size of 2048 bytes for read and write functions
  - Provides a large read request size of 4096 bytes
- The adapter is compatible with 4, 8, and 16 Gb Fibre Channel interface with the following characteristics:
  - Provides for automatic negotiation between 4 Gb, 8 Gb, or 16 Gb link attachments
  - Provides support for the following Fibre Channel topologies: point-to-point (16 Gb only) and fabric
  - Provides support for Fibre Channel class 3
  - Provides a maximum Fibre Channel throughput that is achieved by using full duplex hardware support
- The adapter provides an end-to-end data path parity and CRC protection, including internal data path random-access memory (RAM)
- Provides architectural support for multiple upper layer protocols
- Provides comprehensive virtualization capabilities with support for N_Port ID Virtualization (NPIV) and virtual fabric (VF)
- Provides support for message signaled interrupts extended (MSI-X)
- Provides support for 255 VFs and 1024 MSi-X

- Provides an internal, high-speed static random-access memory (SRAM) memory
- Provides ECC protection of local memory that includes single-bit correction and double-bit protection
- Provides an embedded shortwave optical connection with diagnostics capability
- Provides support for an on-board context management by firmware:
  - Up to 8192 FC port logins
  - I/O multiplexing down to the Fibre Channel frame level
- Provides data buffers capable of supporting 64+ buffer-to-buffer (BB) credits per port for shortwave applications
- Provides link management and recovery that is handled by firmware
- Provides on-board diagnostic capability accessible by an optional connection
- Provides a performance up to 16 Gbps full duplex

Cables:

Cables are the responsibility of the customer. Use multimode fibre optic cables with short-wave lasers that adhere to the following specifications:

- OM4 - multimode 50/125 micron fibre, 4700 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 400m o 8GFC: 0.5m - 190m o 16GFC: 0.5m - 125m o 32GFC: 0.5m - 100m
- OM3 - multimode 50/125 micron fibre, 2000 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 380m o 8GFC: 0.5m - 150m o 16GFC: 0.5m - 100m o 32GFC: 0.5m - 70m
- OM2 - multimode 50/125 micron fibre, 500 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 150m o 8GFC: 0.5m - 50m o 16GFC: 0.5m - 35m o 32GFC: 0.5m - 20m
- OM1 - multimode 62.5/125 micron fibre, 200 MHz*km bandwidth with LC connectors o 4GFC: 0.5m - 70m o 8GFC: 0.5m - 21m o 16GFC: 0.5m - 15m o 32GFC: N/A

Note: The H/W cannot detect what length and type of cable is installed. The link will auto-negotiate to the speed reported during negotiation by the Target. The user must manually set the maximum negotiation speed. If too high of speed is selected, bit errors may occur.
See also optional wrap plug feature #ECW0 which is: a) Required to run some diagnostic procedures and b) In some cases may speed system boot when placed in empty ports as well as avoid useless messages pointing to a planned empty port.

- Attributes provided: 2-port 16Gb Optical FC
- Attributes required: Low profile PCIe Gen3 slot
- Minimum required: 0
- Maximum allowed: 32 (Initial order maximum: 32)
- OS level required:
  - AIX supported
  - IBM i supported
  - Linux supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EP20) Power Enterprise Pools 2.0 Enablement

Power Enterprise Pools 2.0 Enablement

- Attributes provided: Power Enterprise Pools 2.0
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 0)

- OS level required:
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply
  Note: Feature EP20 not orderable in China

(#EP2X) Lab Services Private Cloud Capacity Assessment

Private Cloud Capacity Assessment from Lab Services as part of the Power to Cloud Rewards program, when the engagement is pre-selected via feature #EP2X on one or more qualifying Power E1080 system

- Attributes provided: Private Cloud Capacity Assessment from Lab Services
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
- Initial Order/MES/Both/Supported: Initial
- CSU: Yes
- Return parts MES: No

(#EP2Y) 1 core Mobile Processor Activation

Mobile processor activation for HEX

- Attributes provided: Mobile processor activation
- Attributes required: Inactive Mobile processor cores
- Minimum required: 0
- Maximum allowed: 224 (Initial order maximum: 224)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No

(#EPA0) Deactivation of LPM (Live Partition Mobility)

This feature codes provides firmware commands to deactivate LPM.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EPDC) 1 core Base Processor Activation (Pools 2.0) for EDP2 any OS

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP2 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP2

- Attributes required: #EDP2 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 160 (Initial order maximum: 160)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EPDC not orderable in China

(#EPDD) 1 core Base Processor Activation (Pools 2.0) for EDP3 any OS

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP3 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP3
- Attributes required: #EDP3 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EPDD not orderable in China

(#EPDS) 1 core Base Processor Activation (Pools 2.0) for EDP4 any OS

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP4 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP4
- Attributes required: #EDP4 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EPDS not orderable in China

(#EPDU) 1 core Base Processor Activation (Pools 2.0) for EDP2 Linux only

Each occurrence of this feature will permanently activate one Base Linux only processor core on Processor Card #EDP2 for Pools 2.0

- Attributes provided: One Base Linux only processor core activation (Pools 2.0) for #EDP2
- Attributes required: #EDP2 with inactive Linux only processor cores
- Minimum required: 0
- Maximum allowed: 160 (Initial order maximum: 160)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes

  Return parts MES: No

- Note: Feature EPDU not orderable in China

(#EPDW) 1 core Base Processor Activation (Pools 2.0) for EDP3 Linux only

Each occurrence of this feature will permanently activate one Base Linux only processor core on Processor Card #EDP3 for Pools 2.0

- Attributes provided: One Base Linux only processor core activation (Pools 2.0) for #EDP3
- Attributes required: #EDP3 with inactive Linux only processor cores
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EPDW not orderable in China

(#EPDX) 1 core Base Processor Activation (Pools 2.0) for EDP4 Linux only

Each occurrence of this feature will permanently activate one Base Linux only processor core on Processor Card #EDP4 for Pools 2.0

- Attributes provided: One Base Linux only processor core activation (Pools 2.0) for #EDP4
- Attributes required: #EDP4 with inactive Linux only processor cores
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature EPDX not orderable in China

(#EPE0) Static to Mobile Processor Auto Conversion

Feature is used to order automatic conversion of a static action to a mobile activation without IBM intervention. The term "static activations" also includes Mobile Enabled activations which are technically static. Firmware 840 or later is a pre-requisite. Use of this feature on servers with firmware 830 or earlier will require IBM intervention through the pcod@us.ibm.com

- Attributes provided: Automated conversion
- Attributes required: Firmware 840
- Minimum required: 0
- Maximum allowed: 232 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No

(#EPS0) 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Static)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP2 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP2
- Attributes required: #EDP2 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 160 (Initial order maximum: 160)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPS0 not orderable in China

(#EPS1) 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Static)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP3 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP3
- Attributes required: #EDP3 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPS1 not orderable in China

(#EPS2) 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Static)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP4 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP4
- Attributes required: #EDP4 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPS2 not orderable in China

(#EPS5) 1 core Base Proc Act (Pools 2.0) for #EDP2 Linux (from Static)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP2 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP2
- Attributes required: #EDP2 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 160 (Initial order maximum: 160)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
  CSU: Yes

-
- Return parts MES: No
  Note: Feature EPS5 not orderable in China

(#EPS6) 1 core Base Proc Act (Pools 2.0) for #EDP3 Linux (from Static)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP3 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP3
- Attributes required: #EDP3 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPS6 not orderable in China

(#EPS7) 1 core Base Proc Act (Pools 2.0) for #EDP4 Linux (from Static)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP4 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP4
- Attributes required: #EDP4 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPS7 not orderable in China

(#EPSK) 1 core Base Proc Act (Pools 2.0) for #EDP2 any OS (from Mobile Prev)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP2 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP2
- Attributes required: #EDP2 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 160 (Initial order maximum: 160)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPSK not orderable in China

(#EPSL) 1 core Base Proc Act (Pools 2.0) for #EDP3 any OS (from Mobile Prev)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP3 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP3
- Attributes required: #EDP3 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 192 (Initial order maximum: 192)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPSL not orderable in China

(#EPSM) 1 core Base Proc Act (Pools 2.0) for #EDP4 any OS (from Mobile Prev)

Each occurrence of this feature will permanently activate one Base processor core on Processor Card #EDP4 for Pools 2.0

- Attributes provided: One Base processor core activation (Pools 2.0) for #EDP4
- Attributes required: #EDP4 with inactive processor cores
- Minimum required: 0
- Maximum allowed: 240 (Initial order maximum: 240)
- OS level required: None
- Initial Order/MES/Both/Supported: MES
- CSU: Yes
- Return parts MES: No
  Note: Feature EPSM not orderable in China

(#EPTH) Horizontal PDU Mounting Hardware

This feature ships the hardware required to properly horizontally mount one #EPTG/EPTJ, #EPTK/EPTL, #EPTM/EPTN or #EPTP/EPTQ PDU in a 1U 19-inch rack. A 1U blank panel for the front of the rack for air-flow control is included.

Without this hardware, the PDU can be mounted vertically in the rack's side pockets, but can only be poorly mounted horizontally. The front end of the PDU will be firmly attached to the rear of the rack. But the front of the PDU will be unsupported toward the middle of the rack. Without this hardware, the unsupported end of the PDU will rest on the hardware mounted immediately below it. If that underlying hardware is removed from the rack there is no support for the PDU.

Important Note: This feature code is typically used for an MES order and not for an original order of a new rack with #EPTn PDUs. As part of factory integration, IBM Manufacturing automatically adds this hardware without a feature code and at no additional charge when its #EPTn PDU placement logic calls for horizontally mounted PDUs. Use this feature code when (1) converting an existing vertically mounted #EPTn PDU to horizontal mounting or (2) separately ordering a #EPTn PDU for horizontal field installation.

- Attributes provided: mounting hardware
- Attributes required: High Function PDU (#EPT*) and space in 19-inch rack
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EPTJ) High Function 9xC19 PDU: Switched, Monitoring

(No longer available as of April 24, 2020)

This is an intelligent, switched 200-240 volt AC Power Distribution Unit (PDU) with nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the nine C19 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Depending on country wiring standards the PDU is single-phase or three-phase wye. See three-phase #EPTK/EPTL for countries which do not use wye wiring.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C20 plug connect to C19 PDU receptacles and are ordered separately. One country-specific wall line cord is also ordered separately and attaches to a UTG524-7 connector on the front of the PDU. Supported line cords include features #6489, #6491, #6492, #6653, #6654, #6655, #6656, #6657, #6658, #6667, #ELC1 or #ELC2.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

There are also three C13 receptacles on the rear of the PDU positioned toward the middle of the rack. These are generally not easily accessed and therefore IBM does not generally recommend their use.

#EPTG and #EPTJ are identical PDUs. Up to one lower price #EPTG can be ordered with a new 7014-T42/T00 rack in place of a no-charge #9188 PDU.

For comparison, this is most similar to the earlier generation #7189 PDU.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Nine C19 PDU - switched, power monitoring
- Attributes required: PDU wall line cord & space in 19-inch rack
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EPTL) High Function 9xC19 PDU 3-Phase

Switched, Monitoring

This is an intelligent, switched 208 volt 3-phase AC Power Distribution Unit (PDU) with nine C19 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the nine C19 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C20 plug connect to C19 PDU receptacles and are ordered separately. One wall line cord is provided with the PDU (no separate feature code) and has a IEC60309 60A plug (3P+G). The PDU supports up to 48 amps.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

There are also three C13 receptacles on the rear of the PDU positioned toward the middle of the rack. These are generally not easily accessed and therefore IBM does not generally recommend their use.

#EPTK and #EPTL are identical PDUs. Up to one lower price #EPTK can be ordered with a new 7965-S42 rack in place of a no-charge #9188 PDU.

For comparison, this is most similar to the earlier generation #7196 PDU.

Not orderable in China and Hong Kong.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Nine C19 PDU - switched, power monitoring
- Attributes required: space in rack, 3-phase 208V AC delta electrical service
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EPTN) High Function 12xC13 PDU: Switched, Monitoring

(No longer available as of April 24, 2020)

This is an intelligent, switched 200-240 volt AC Power Distribution Unit (PDU) with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the twelve C13 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker. Depending on country wiring standards the PDU is single- phase or three-phase wye. See three-phase #EPTK/EPTL for countries which do not use wye wiring.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C14 plug connect to C13 PDU receptacles and are ordered separately. One country-specific wall line cord is also ordered separately and attaches to a UTG524-7 connector on the front of the PDU. Supported line cords include features #6489, #6491, #6492, #6653, #6654, #6655, #6656, #6657, #6658, #6667, #ELC1 or #ELC2.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

#EPTM and #EPTN are identical PDUs. Up to one lower price #EPTM can be ordered with a new 7014-T42/T00 rack in place of a no-charge #9188 PDU.

For comparison, this is most similar to the earlier generation #7109 PDU.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Twelve C13 PDU - switched, power monitoring
- Attributes required: PDU wall line cord & space in 19-inch rack
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EPTQ) High Function 12xC13 PDU 3-Phase

Switched, Monitoring

This is an intelligent, switched 208 volt 3-phase AC Power Distribution Unit (PDU) with twelve C13 receptacles on the front of the PDU. The PDU is mounted on the rear of the rack making the twelve C13 receptacles easily accessible. Each receptacle has a 20 amp circuit breaker.

The PDU can be mounted vertically in rack side pockets or it can be mounted horizontally. If mounted horizontally, it uses 1 EIA (1U) of rack space. See feature #EPTH for horizontal mounting hardware.

Device power cords with a C14 plug connect to C13 PDU receptacles and are ordered separately. One wall line cord is provided with the PDU (no separate feature code) and has a IEC60309 60A plug (3P+G). The PDU supports up to 48 amps.

Two RJ45 ports on the front of the PDU enable the client to monitor each receptacle's electrical power usage and to remotely switch any receptacle on or off. The PDU is shipped with a generic PDU password and IBM strongly urges clients to change it upon installation.

#EPTP and #EPTQ are identical PDUs. Up to one lower price #EPTP can be ordered with a new 7965-S42 rack in place of a no-charge #9188 PDU.

For comparison, this is most similar to the earlier generation #7196 PDU, but offers C13 receptacles.

Not orderable in China and Hong Kong.

Limitation: Some configurations of the Elastic Storage Server (ESS) are delivered with a Intelligent PDU. At this time, the intelligent management capabilities of this PDU are not configured or used by the ESS system. If the ESS Customer would like to use this capability, it is the Customers responsibility to configure this PDU. In any case the ethernet port on the Intelligent PDU must not be connected to the ESS Management switch.

- Attributes provided: Twelve C13 PDU - switched, power monitoring
- Attributes required: space in rack, 3-phase 208V AC delta electrical service
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 0)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EQ77) Qty 150 of #6577

This feature provides QTY 150 of #6577. FC 6577 is mandatory on factory rack integrated (4651-4666) orders. Feature is not valid on initial order with non-factory integrated feature 4650.

- Attributes provided: Power cable
- Attributes required: At least one Rack and the absence of #4650.
- Minimum required: 0
- Maximum allowed: 2 (Initial order maximum: 2)
- OS level required: None
- Initial Order/MES/Both/Supported: Supported
- CSU: N/A
- Return parts MES: No

(#ER16) Indicator, reserve 5 EIA rack space

Virtual enclosure indicator to reserve 5U adjacent rack space for system node.

- Attributes provided: Reserve 5U rack space for system node.
- Attributes required: 5U of rack space

- Minimum required: 0
- Maximum allowed: 3 (Initial order maximum: 3)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#ER1A) Specify Reserve 4 EIA Rack Space for PCIe3 Expansion Drawer

Virtual enclosure indicator to reserve 4U adjacent rack space for I/O expansion drawer.

- Attributes provided: Reserve 4U rack space for PCIe Gen3 I/O Expansion Drawer
- Attributes required: 4U rack space adjacent to installed I/O expansion drawer
- Minimum required: 0
- Maximum allowed: 4 (Initial order maximum: 4)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: Does not apply

(#ER21) Field Integration of Rack and Server

IBM Manufacturing will test the server in a 7964-S42 rack then remove the server from the rack and ship the server in separate packages without a rack for field integration.

- Attributes provided: Customer provided rack.
- Attributes required: IBM manufacturing testing and then shipment of server without a rack.
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#ES94) 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ES94 and #ES95 are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ES94 indicates usage by AIX, Linux or VIOS. Feature ES95 indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i not supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ES95) 387GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ES94 and #ES95 are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ES94 indicates usage by AIX, Linux or VIOS. Feature ES95 indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESB2) 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 5xx

(528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESB2 CCIN is 5B16. Different feature codes to help the IBM configuration tools understand how the SSD is used. ESB2 indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 5xx SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - AIX - supported
  - IBM i not supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESB6) 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 5xx (528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESGZ CCIN is 5B17. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESGZ indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access

the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 5xx SSD
- Attributes required: SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - AIX - supported
  - IBM i not supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESBA) 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESBA and #ESBB are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESBA indicates usage by AIX, Linux or VIOS. Feature ESBB indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - AIX - supported
  - IBM i not supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESBB) 387GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESBA and #ESBB are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESBA indicates usage by AIX, Linux or VIOS. Feature ESBB indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. The 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESBG) 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESBG and #ESBH are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESBG indicates usage by AIX, Linux or VIOS. Feature ESBH indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Attributes provided: one 775 GB SFF-2 4k SSD

-
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - AIX - supported
  - IBM i not supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESBH) 775GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESBG and #ESBH are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESBG indicates usage by AIX, Linux or VIOS. Feature ESBH indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. The 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESBL) 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESBL and #ESBM are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESBL indicates usage by AIX, Linux or VIOS. Feature ESBM indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - AIX - supported
  - IBM i not supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESBM) 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESBL and #ESBM are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESBL indicates usage by AIX, Linux or VIOS. Feature ESBM indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. The 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)

- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESC0) S&H - No Charge

No charge shipping and handling

- Attributes provided: None
- Attributes required: Sales Preapproval Required
- Minimum required: 0
- Maximum allowed: 3 (Initial order maximum: 1)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: N/A
- Return parts MES: Does not apply

(#ESC9) S&H

Shipping and handling

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 3 (Initial order maximum: 1)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply

(#ESCZ) - iSCSI SAN Load Source Specify for AIX

Indicates that an LAN adapter is being used as the Load Source for the AIX operating system.

- Attributes provided: iSCSI load source placement specify
- Attributes required: LAN adapter
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - AIX supported
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES: No

(#ESEU) 571GB 10K RPM SAS SFF- HDD 4K for IBM i

571 GB 2.5-inch (Small Form Factor (SFF)) 10k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays such as found in EXP24S storage drawer. Disk is formatted for 4224 byte sectors. If reformatted to 4096 byte sectors, its capacity would be 500 GB and it would not have the 4224 byte additional data integrity protection. Reformatting large drives takes significant time.

Features #ESEU and #ESEV are physically identical drives with the same 59D2 CCIN. Different feature codes help the IBM configuration tools understand how the HDD is used. #ESEV indicates usage by AIX, Linux or VIOS. #ESEU indicates usage by IBM i.

Limitations:

- Cannot be combined in the same array as a drive using different sector size
- Physically does not fit in a SFF-1 or SFF-3 bay due to carrier/ tray differences
- Attributes provided: 571GB 10K RPM SFF-2 Disk 4K
- Attributes required: one SFF-2 SAS bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESEV) 600GB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux

600 GB 2.5-inch (Small Form Factor (SFF)) 10k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays such as found in EXP24S storage drawer. IBM Manufacturing may ship formatted with 4224 byte sectors or with 4096 sectors. With 4096 byte sectors the drive's capacity is 600 GB or with 4224 byte sectors the capacity is 571 GB. Using 4224 byte sectors provides additional data integrity protection. Reformatting large drives takes significant time.

Features #ESEU and #ESEV are physically identical drives with the same 59D2 CCIN. Different feature codes help the IBM configuration tools understand how the HDD is used. #ESEV indicates usage by AIX, Linux or VIOS. #ESEU indicates usage by IBM i.

Limitations:

- Cannot be combined in the same array as a drive using different sector size
- Physically does not fit in a SFF-1 or SFF-3 bay due to carrier/ tray differences

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: 571GB 10K RPM SFF-2 Disk 4K
- Attributes required: one SFF-2 SAS bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - AIX - supported
  - IBM i not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESF2) 1.1TB 10K RPM SAS SFF-2 HDD 4K for IBM i

1.14 TB 2.5-inch (Small Form Factor (SFF)) 10k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays such as found in EXP24S storage drawer. Disk is formatted for 4224 byte sectors. If reformatted to 4096 byte sectors, its capacity would be 1.2 TB and it would not have the 4224 byte additional data integrity protection. Reformatting large drives takes significant time.

Features #ESF2 and #ESF3 are physically identical drives with the same 59DA CCIN. Different feature codes help the IBM configuration tools understand how the HDD is used. #ESF3 indicates usage by AIX, Linux or VIOS. #ESF2 indicates usage by IBM i.

Limitations:

- Cannot be combined in the same array as a drive using different sector size
- Physically does not fit in a SFF-1 or SFF-3 bay due to carrier/ tray differences
- Attributes provided: 1.1TB 10K RPM SFF-2 Disk 4K
- Attributes required: one SFF-2 SAS bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESF3) 1.2TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux

1.2 TB 2.5-inch (Small Form Factor (SFF)) 10k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays such as found in EXP24S storage drawer. IBM Manufacturing may ship formatted with 4224 byte sectors or with 4096 sectors. With 4096 byte sectors the drive's capacity is 1.2 TB or with 4224 byte sectors the capacity is 1.14TB. Using 4224 byte sectors provides additional data integrity protection. Reformatting large drives takes significant time.

Features #ESF2 and #ESF3 are physically identical drives with the same 59DA CCIN. Different feature codes help the IBM configuration tools understand how the HDD is used. #ESF3 indicates usage by AIX, Linux or VIOS. #ESF2 indicates usage by IBM i.

Limitations:

- Cannot be combined in the same array as a drive using different sector size
- Physically does not fit in a SFF-1 or SFF-3 bay due to carrier/ tray differences

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: 1.2TB 10K RPM SFF-2 Disk 4K
- Attributes required: one SFF-2 SAS bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - AIX - supported
  - IBM i not supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both

- CSU: Yes
- Return parts MES: No

(#ESFS) 1.7TB 10K RPM SAS SFF-2 HDD 4K for IBM i

1.71 TB 2.5-inch (Small Form Factor (SFF)) 10k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays such as found in EXP24S storage drawer. Disk is formatted for 4224 byte sectors. If reformatted to 4096 byte sectors, its capacity would be 1.8 TB and it would not have the 4224 byte additional data integrity protection. Reformatting large drives takes significant time.

Features #ESFS and #ESFT are physically identical drives with the same 59DD CCIN. Different feature codes help the IBM configuration tools understand how the HDD is used. #ESFT indicates usage by AIX, Linux or VIOS. #ESFS indicates usage by IBM i.

Limitations:

- Cannot be combined in the same array as a drive using different sector size
- Physically does not fit in a SFF-1 or SFF-3 bay due to carrier/ tray differences
- Attributes provided: 1.7TB 10K RPM SFF-2 Disk 4K
- Attributes required: one SFF-2 SAS bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESFT) 1.8TB 10K RPM SAS SFF-2 HDD 4K for AIX/Linux

1.8 TB 2.5-inch (Small Form Factor (SFF)) 10k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays such as found in EXP24S storage drawer. IBM Manufacturing may ship formatted with 4224 byte sectors or with 4096 sectors. With 4096 byte sectors the drive's capacity is 1.8 TB or with 4224 byte sectors the capacity is 1.71TB. Using 4224 byte sectors provides additional data integrity protection. Reformatting large drives takes significant time.

Features #ESFS and #ESFT are physically identical drives with the same 59DD CCIN. Different feature codes help the IBM configuration tools understand how the HDD is used. #ESFT indicates usage by AIX, Linux or VIOS. #ESFS indicates usage by IBM i.

Limitations:

- Cannot be combined in the same array as a drive using different sector size
- Physically does not fit in a SFF-1 or SFF-3 bay due to carrier/ tray differences

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: 1.8TB 10K RPM SFF-2 Disk 4K
- Attributes required: one SFF-2 SAS bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  AIX - supported

-
- IBM i not supported
- Linux - supported

Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESGV) 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 5xx (528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESGV CCIN is 5B16. Different feature codes to help the IBM configuration tools understand how the SSD is used. ESGV indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 5xx SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i not supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESGZ) 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 5xx (528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESGZ CCIN is 5B17. Different feature codes to help the IBM configuration tools understand how the SSD is

used. Feature ESGZ indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 5xx SSD
- Attributes required: SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i not supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ0) 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.45 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJ0 and #ESJ1 are physically identical drives with the same 5B29 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ0 indicates usage by AIX, Linux or VIOS. #ESJ1 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do

not support 4k drives.
- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter or later.
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i not supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ1) 931GB Mainstream SAS 4k SFF-2 SSD for IBM i

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.45 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and

Features #ESJ0 and #ESJ1 are physically identical drives with the same 5B29 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ0 indicates usage by AIX, Linux or VIOS. #ESJ1 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ2) 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.45 TB 13601
```

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJ2 and #ESJ3 are physically identical drives with the same 5B21 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ2 indicates usage by AIX, Linux or VIOS. #ESJ3 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i not supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ3) 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
```

3.72 TB 6799
7.45 TB 13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJ2 and #ESJ3 are physically identical drives with the same 5B21 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ2 indicates usage by AIX, Linux or VIOS. #ESJ3 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ4) 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.45 TB 13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJ4 and #ESJ5 are physically identical drives with the same 5B2D CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ4 indicates usage by AIX, Linux or VIOS. #ESJ5 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i not supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ5) 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.45 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJ4 and #ESJ5 are physically identical drives with the same 5B2D CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ4 indicates usage by AIX, Linux or VIOS. #ESJ5 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)

  OS level required:

- IBM i supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ6) 7.45TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

7.45 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.45 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJ6 and #ESJ7 are physically identical drives with the same 5B2F CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ6 indicates usage by AIX, Linux or VIOS. #ESJ7 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 7.45 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i not supported
  - AIX supported
  - Linux supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJ7) 7.45TB Mainstream SAS 4k SFF-2 SSD for IBM i

7.45 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.45 TB 13601
```

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJ6 and #ESJ7 are physically identical drives with the same 5B2F CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJ6 indicates usage by AIX, Linux or VIOS. #ESJ7 indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 7.45 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
  OS level required:
  - IBM i 7.5, or later
  - IBM i 7.4 TR5, or later
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJJ) 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.45 TB 13601
```

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJJ and #ESJK are physically identical drives with the same 5B29 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJJ indicates usage by AIX, Linux or VIOS. #ESJK indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter or later.
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - Linux - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJK) 931GB Mainstream SAS 4k SFF-2 SSD for IBM i

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.45 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJJ and #ESJK are physically identical drives with the same 5B29 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJJ indicates usage by AIX, Linux or VIOS. #ESJK indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.

- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter or later.
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJL) 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.45 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJL and #ESJM are physically identical drives with the same 5B21 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJL indicates usage by AIX, Linux or VIOS. #ESJM indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJM) 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.45 TB 13601
```

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJL and #ESJM are physically identical drives with the same 5B21 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJL indicates usage by AIX, Linux or VIOS. #ESJM indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJN) 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.45 TB 13601
```

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions

section or IBM Documentation for additional detail.

Features #ESJN and #ESJP are physically identical drives with the same 5B2D CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJN indicates usage by AIX, Linux or VIOS. #ESJP indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJP) 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.45 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJN and #ESJP are physically identical drives with the same 5B2D CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJN indicates usage by AIX, Linux or VIOS. #ESJP indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's

SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.

- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESJQ) 7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

7.44 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.44 TB	13601

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJQ and #ESJR are physically identical drives with the same 5B2F CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJQ indicates usage by AIX, Linux or VIOS. #ESJR indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 7.44 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes

  Return parts MES: No

-

(#ESJR) 7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i

7.44 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Depending on the nature of the workload, the lifetime TBW may be somewhat larger. To read the warranty and maintenance applicable to mainstream devices on POWER8 and POWER9 servers, see the Terms and Conditions section or IBM Documentation for additional detail.

Features #ESJQ and #ESJR are physically identical drives with the same 5B2F CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESJQ indicates usage by AIX, Linux or VIOS. #ESJR indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER8/POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 7.44 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
  OS level required:
  - IBM i 7.5, or later
  - IBM i 7.4 TR5, or later
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESK1) 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 5xx (528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESK1 CCIN is 5B16. Different feature codes to help the IBM configuration tools understand how the SSD is

used. ESK1 indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives can not be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 5xx SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESK3) 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 5xx (528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESK3 CCIN is 5B17. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESK3 indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives can not be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 5xx SSD
- Attributes required: SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - not supported

    - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESK8) 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESK8 and #ESK9 are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESK8 indicates usage by AIX, Linux or VIOS. Feature ESK9 indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives can not be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
    - IBM i - not supported
    - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESK9) 387GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESK8 and #ESK9 are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESK8 indicates usage by AIX, Linux or VIOS. Feature ESK9 indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. The 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKC) 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESKC and #ESKD are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESKC indicates usage by AIX, Linux or VIOS. Feature ESKD indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives can not be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKD) 775GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESKC and #ESKD are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESKC indicates usage by AIX, Linux or VIOS. Feature ESKD indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives can not be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKG) 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESKG and #ESKH are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESKG indicates usage by AIX, Linux or VIOS. Feature ESKH indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives can not be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access

the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKH) 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ESKG and #ESKH are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESKG indicates usage by AIX, Linux or VIOS. Feature ESKH indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 system units (SFF-3) or in older SFF-1 SAS bays. The 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKK) 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.44 TB	13601

Features #ESKK and #ESKM are physically identical drives with the same 5B29 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKK indicates usage by AIX, Linux or VIOS. #ESKM indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKM) 931GB Mainstream SAS 4k SFF-2 SSD for IBM i

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.44 TB	13601

Features #ESKK and #ESKM are physically identical drives with the same 5B29 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKK indicates usage by AIX, Linux or VIOS. #ESKM indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter or later.
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKP) 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.44 TB	13601

Features #ESKP and #ESKR are physically identical drives with the same 5B21 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKP indicates usage by AIX, Linux or VIOS. #ESKR indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - not supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKR) 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Features #ESKP and #ESKR are physically identical drives with the same 5B21 CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKP indicates usage by AIX, Linux or VIOS. #ESKR indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKT) 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Features #ESKT and #ESKV are physically identical drives with the same 5B2D CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKT indicates usage by AIX, Linux or VIOS. #ESKV indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx

byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKV) 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Features #ESKT and #ESKV are physically identical drives with the same 5B2D CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKT indicates usage by AIX, Linux or VIOS. #ESKV indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - supported
- Initial Order/MES/Both/Supported: Both

- CSU: Yes
- Return parts MES: No

(#ESKX) 7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

7.44 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.44 TB	13601

Features #ESKX and #ESKZ are physically identical drives with the same 5B2F CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKX indicates usage by AIX, Linux or VIOS. #ESKZ indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 7.44 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - not supported
  - Linux - supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESKZ) 7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i

7.44 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799

7.44 TB 13601

Features #ESKX and #ESKZ are physically identical drives with the same 5B2F CCIN. Different feature codes help the IBM configuration tools understand how the SSD is used. #ESKX indicates usage by AIX, Linux or VIOS. #ESKZ indicates usage by IBM i.
Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/ 528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive can not be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 7.44 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
  OS level required:
  - IBM i 7.5, or later
  - IBM i 7.4 TR5, or later supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESLA) Specify AC Power Supply for EXP12SX/EXP24SX Storage Enclosure

No-charge specify for AC power supply for an EXP12SX or EXP24SX SAS Storage Enclosure.

The power supply has a 320-C14 inlet electrical connection for a separately ordered power cord. It is rated 800 Watts Output Power and 100 - 240 VAC (RMS) input voltage.

- Attributes provided: communicates to IBM Manufacturing an AC Power Supply is to be used
- Attributes required: EXP12SX or EXP24SX Storage Enclosure and AC power cord
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 0)
- OS level required: N/A
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESLB) ESBB Load Source Specify (387GB SSD SFF-2)

This specify code indicates that a #ESBB Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESBB
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature ESBB
- Initial Order/MES/Both/Supported: Supported

- CSU: Yes
- Return parts MES: No

(#ESLH) ESBH Load Source Specify (775GB SSD SFF-2)

This specify code indicates that a #ESBH Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESBH
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature ESBH
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESLM) ESBM Load Source Specify (1.55TB SSD SFF-2)

This specify code indicates that a #ESBM Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESBM
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature ESBM
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESLS) 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-swap Hard Disk Drives (HDD) or Solid State Drives (SSD) in only 2 EIA of space in a 19-inch rack. The EXP24SX SFF bays use SFF gen2 (SFF-2) carriers/ trays identical to the carrier/trays in the previous EXP24S Drawer .

With AIX/Linux/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 one set of 24 bays (mode 1) is supported. The mode setting can be changed in the field using software commands along with a specificly documented procedure. IMPORTANT NOTE: when changing mode, it is very important that you follow the documented procedures and that there is no data on the drives before the change. 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 re-configuration work.

The EXP24SX has redundant SAS paths to all drives via two redundant Enclosure Services Modules (ESMs). Four mini-SAS HD narrow ports are attached to PCIe Gen3 SAS adapters such as the #EJ0J/EJ0M or #EJ0L or #EJ14, or attached to an imbedded SAS controller in a POWER8/POWER9 Scale-out server such as the Power S814, S822, S824, S914, S922, S924, H922, or H924. Attachment between the SAS controller and the storage enclosure SAS ports is via the appropriate SAS YO12 or X12 cables. The PCIe Gen3 SAS adapters support 6Gb throughput. The EXP24SX has been designed to support up to 12Gb throughput if future SAS adapters support that capability.

The EXP24SX uses redundant power supplies and two power cords. Order two feature #ESLA for AC power supplies. The enclosure is shipped with adjustable depth rails and can accommodate rack depths from 59.5 - 75 cm (23.4 - 29.5 inches). Slot filler panels are provided for empty bays when initially shipped from IBM.

See also the 12-bay Large Form Factor (LFF) EXP12SX SAS Storage Enclosure (feature #ESLL) for higher capacity drives with lower performance.

Attributes provided: 24 SFF-2 SAS bays in a 2U enclosure

-
- Attributes required: PCIe Gen3 SAS adapter/controller; Power System (at least POWER8/ POWER9 generation); 2U 19-inch rack space; Appropriate SAS cables
- Minimum required: 0
- Maximum allowed: 168 (Initial order maximum: 168)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: VIOS supported.

(#ESMB) - 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.44 TB	13601

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESMD) - 931GB Mainstream SAS 4k SFF-2 SSD for IBM i

931 GB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual

performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 931 GB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter or later.
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESMF) - 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:

- IBM i - Not Supported
- AIX - Supported
- Linux - Supported

Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESMH) - 1.86TB Mainstream SAS 4k SFF-2 SSD for IBM i

1.86 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399
3.72 TB	6799
7.44 TB	13601

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 1.86 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESMK) - 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

Drive Capacity	Total Bytes Written (TBW) in (TB)
931 GB	1700
1.86 TB	3399

```
3.72 TB 6799
7.44 TB 13601
```

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESMS) - 3.72TB Mainstream SAS 4k SFF-2 SSD for IBM i

3.72 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Limitation:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 3.72 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both

  CSU: Yes

-
- Return parts MES: No

(#ESMV) - 7.44TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux

7.44 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do not support 4k drives.
- Attributes provided: 7.44 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESMX) - 7.44TB Mainstream SAS 4k SFF-2 SSD for IBM i

7.44 TB SAS 2.5-inch (SFF) Mainstream solid state drive (SSD) formatted in 4224 byte sectors (4k). The drive is mounted on a 2.5-inch SFF-2 carrier/tray to fit an expansion drawer such as the EXP24SX. Drive is rated at 6Gb/s though actual performance is dependent upon multiple factors.

DWPD (Drive Write Per Day) rating is 1 calculated over a 5 year period. See the following table for the approximate lifetime Total Bytes Written (TBW) supported by each drive capacity:

```
Drive Capacity Total Bytes Written (TBW) in (TB)
931 GB 1700
1.86 TB 3399
3.72 TB 6799
7.44 TB 13601
```

Limitations:

- The drive is mounted on a SFF-2 carrier/tray and does not physically fit into a POWER9 or POWER10 system unit's SFF-3 bays. JBOD formatting of 4096 byte sectors is not tested or supported. Drive cannot be reformatted to 5xx byte sectors. 5xx and 4k drives cannot be mixed in the same array. Older SAS adapters such as PCIe2 or earlier do

not support 4k drives.

- Attributes provided: 7.44 TB Mainstream SSD
- Attributes required: Open SFF-2 SAS bay controlled by PCIe Gen3 PCIe SAS adapter, or later
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
  OS level required:
  - IBM i 7.5, or later
  - IBM i 7.4 TR5, or later Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESNA) 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESNA and #ESNB are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESNA indicates usage by AIX, Linux or VIOS. Feature ESNB indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESNB) 775GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESNA and #ESNB are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESNA indicates usage by AIX, Linux or VIOS. Feature ESNB indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESNE) 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESNE and #ESNF are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESNE indicates usage by AIX, Linux or VIOS. Feature ESNF indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not

supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESNF) 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24S (#5887) or EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM solid state device failures will be replaced during the 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.

Features #ESNE and #ESNF are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ESNE indicates usage by AIX, Linux or VIOS. Feature ESNF indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER8/POWER9 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#ESNL) 283GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)

283 GB 2.5-inch (Small Form Factor (SFF)) 15k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays. Disk is formatted for 4224 byte sectors. If reformatted to 4096 byte sectors, capacity would be 300 GB. Drive includes enhanced caching capability with 256MB DRAM which may improve performance somewhat compared to previous disk drives without the cache.

CCIN is 5B43.

Limitations:

- Cannot be combined in the same array as a drive of the same capacity, but using different sector size. However, can be combined with non-cached drives of same capacity in the same array.
- This HDD feature is not supported in the EXP24S drawer (#5887 or #EL1S).

- Attributes provided: 283 GB Disk Drive - SFF-2
- Attributes required: one SFF-2 drive bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESNM) 300GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)

300 GB 2.5-inch (Small Form Factor (SFF)) 15k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays. Disk is formatted for 4096 byte sectors. If reformatted to 4224 byte sectors, capacity would be 283 GB. Drive includes enhanced caching capability with 256MB DRAM which may improve performance somewhat compared to previous disk drives without the cache.

CCIN is 5B43.

Limitations:

- Cannot be combined in the same array as a drive of the same capacity, but using different sector size. However, can be combined with non-cached drives of same capacity in the same array.
- This HDD feature is not supported in the EXP24S drawer (#5887 or #EL1S).
- Attributes provided: 300 GB Disk Drive - SFF-2
- Attributes required: one SFF-2 drive bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux -Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESNQ) 571GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (IBM i)

571 GB 2.5-inch (Small Form Factor (SFF)) 15k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays. Disk is formatted for 4224 byte sectors. If reformatted to 4096 byte sectors, capacity would be 600 GB. Drive includes enhanced caching capability with 256MB DRAM which may improve performance somewhat compared to previous disk drives without the cache.

CCIN is 5B47.

Limitations:

- Cannot be combined in the same array as a drive of the same capacity, but using different sector size. However, can be combined with non-cached drives of same capacity in the same array.
- This HDD feature is not supported in the EXP24S drawer (#5887 or #EL1S).
- Attributes provided: 571 GB Disk Drive - SFF-2
- Attributes required: one SFF-2 drive bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESNR) 600GB 15K RPM SAS SFF-2 4k Block Cached Disk Drive (AIX/Linux)

600 GB 2.5-inch (Small Form Factor (SFF)) 15k rpm SAS disk drive on Gen-2 carrier/tray. Supported in SFF-2 SAS bays. Disk is formatted for 4096 byte sectors. If reformatted to 4224 byte sectors, capacity would be 571 GB. Drive includes enhanced caching capability with 256MB DRAM which may improve performance somewhat compared to previous disk drives without the cache.

CCIN is 5B47.

Limitations:

- Cannot be combined in the same array as a drive of the same capacity, but using different sector size. However, can be combined with non-cached drives of same capacity in the same array.
- This HDD feature is not supported in the EXP24S drawer (#5887 or #EL1S).
- Attributes provided: 600 GB Disk Drive - SFF-2
- Attributes required: one SFF-2 drive bay
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 4032)
- OS level required:
  - IBM i - Not supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ESWK) AIX Update Access Key (UAK)

With the introduction of the Power10 server, IBM expands upon the use of the update access key (UAK) with the addition of an AIX Software Maintenance (SWMA) UAK. This functionality provides proactive notification of AIX SWMA expirations to ensure continued and uninterrupted software support. AIX SWMA UAKs do not limit the operability of or capability to update AIX. For additional information see announcement letters:

- For Japan: JP21-0387
- For Asia Pacific: AP21-0322
- For United States: 221-331
- For Latin America: LP21-0395
- For Canada: A21-050
- For Europe Middle East and Africa: ZP21-0403
- Attributes provided: AIX software maintenance agreements (SWMA) validation
- Attributes required: AIX OS on server
- Minimum required: 0
- Maximum allowed: 999 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Initial
- CSU: Yes
- Return parts MES: No

(#ETK1) - 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 5xx (528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETK1 CCIN is 5B16. Different feature codes to help the IBM configuration tools understand how the SSD is used. ETK1 indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access

the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 5xx SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ETK3) - 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 5xx (528) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETK3 CCIN is 5B17. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ETK3 indicates usage by AIX, Linux or VIOS.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 512 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 5xx SSD
- Attributes required: SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ETK8) - 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETK8 and #ETK9 are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ETK8 indicates usage by AIX, Linux or VIOS. Feature ETK9 indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ETK9) - 387GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 387 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETK8 and #ETK9 are physically identical drives with the same CCIN of 5B10. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ETK8 indicates usage by AIX, Linux or VIOS. Feature ETK9 indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in

older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 387 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ETKC) - 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETKC and #ETKD are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ETKC indicates usage by AIX, Linux or VIOS. Feature ETKD indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported

- Linux - Supported

Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ETKD) - 775GB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 775 GB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETKC and #ETKD are physically identical drives with the same CCIN of 5B11. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ETKC indicates usage by AIX, Linux or VIOS. Feature ETKD indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 775 GB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported

- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ETKG) - 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETKG and #ETKH are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ETKG indicates usage by AIX, Linux or VIOS. Feature ETKH indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD
- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Not Supported
  - AIX - Supported
  - Linux - Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#ETKH) - 1.55TB Enterprise SAS 4k SFF-2 SSD for IBM i

This SFF (2.5") Enterprise SAS Solid State Drive (SSD) is shipped from IBM with 1.55 TB capacity formatted with 4k (4224) byte sectors. The drive is supported in SFF-2 SAS bays in the EXP24SX (#ESLS) I/O drawer. This drive is rated at 10 DWPD (Drive Writes Per Day) calculated over a 5-year period.

IBM SDD 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.

Features #ETKG and #ETKH are physically identical drives with the same CCIN of 5B12. Different feature codes to help the IBM configuration tools understand how the SSD is used. Feature ETKG indicates usage by AIX, Linux or VIOS. Feature ETKH indicates usage by IBM i.

Limitations:

- Due to physical differences in the carrier/tray on which the drive is mounted, the SSD only fits in SFF-2 bays such as found in EXP24SX storage enclosure and cannot be used in POWER9 or POWER10 system units (SFF-3) or in older SFF-1 SAS bays. Also the drive was not tested with 4096 byte sectors and thus JBOD mode for AIX/Linux is not supported. 4k drives cannot be reformatted to 5xx drives (or vice versa).

Note: As part of a larger industry transition, IBM Power Systems is in the process of gradually shifting away from offering 5xx (512/528) byte sector SAS drives (HDD or SSD) to 4K byte sector drives. Most AIX/Linux applications access SAS drives through a file system and experience no effect of this sector size transition. In much less common cases, an AIX/Linux application may access the disk drive directly through the raw device interface and may still require only 5xx byte sectors drives. If there are 4k usage concerns, check with the application provider to ensure compatibility with 4K byte sector drives.

- Attributes provided: one 1.55 TB SFF-2 4k SSD

- Attributes required: one SFF-2 SAS bay, PCIe3 SAS controller
- Minimum required: 0
- Maximum allowed: 4032 (Initial order maximum: 999)
- OS level required:
  - IBM i - Supported
  - AIX - Not Supported
  - Linux - Not Supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EU01) 1TB Removable Disk Drive Cartridge

1TB Removable Disk Drive Cartridge (#EU01) provides a RDX disk drive in a rugged cartridge to be used in an RDX Internal and External Docking Station such as the (#1103, #1104 or #1123, #EU03, #EU04 ,#EU23, or #EU07 1TB is uncompressed. docking station. 1TB is uncompressed. Compression/decompression is provided by the operating system, not the drive itself. Feature EU01 is not entitled under the IBM Maintenance Agreement, if one is purchased.

- Attributes provided: 1TB RDX rugged disk cartridge
- Attributes required: None.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: see RDX docking station
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EU29) Order Placed Indicator

This feature is use to identify ORDER PLACED for administrative purposes within manufacturing to Facilitate processing.

- Attributes provided: ORDER PLACED INDICATOR FOR ADMINISTRATIVE TRACKING
- Attributes required: None.
- Minimum required: 0
- Maximum allowed: 9999 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: Feature orderable in Algeria, Comoros or Tunisia only

(#EU2T) 2TB Removable Disk Drive Cartridge (RDX)

The 2.0TB Removable Disk Drive Cartridge provides a RDX disk drive in a rugged cartridge to be used in an RDX Internal and External Docking Station such as (#1103, #1104 or #1123, #EU03, #EU04 ,#EU23, or #EU07. 2.0TB is uncompressed. Compression/ decompression is provided by the operating system, not the drive itself. Feature EU2T is not entitled under the IBM Maintenance Agreement, if one is purchased.

- Attributes provided: 2.0TB RDX rugged disk cartridge
- Attributes required: None.
- Minimum required: 0

- Maximum allowed: 9999 (Initial order maximum: 250)
- OS level required: See docking station for OS requirements
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EU41) ESJ1 Load Source Specify (931GB SSD SFF-2)

This specify code indicates that a #ESJ1 Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESJ1
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature #ESJ1
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EU43) ESJ3 Load Source Specify (1.86TB SSD SFF-2)

This specify code indicates that a #ESJ3 Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESJ3
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature #ESJ3
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EU45) ESJ5 Load Source Specify (3.72TB SSD SFF-2)

This specify code indicates that a #ESJ5 Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESJ5
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature #ESJ5
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EU47) ESJ7 Load Source Specify (7.45TB SSD SFF-2)

This specify code indicates that a #ESJ7 Solid State Drive is being used as the Load Source.

- Attributes provided: Communicate location of load source to IBM Manufacturing
- Attributes required: Feature ESJ7

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: see feature #ESJ7
- Initial Order/MES/Both/Supported: Supported
- CSU: Yes
- Return parts MES: No

(#EUA4) RDX USB External Docking Station

USB External Docking Station which accommodates RDX removable disk cartridge of any capacity. The disk is in a protective rugged cartridge enclosure that plug into the docking station. The docking station holds one removable rugged disk drive/cartridge at a time. The rugged removable disk cartridge and docking station can be used similar to a tape drive. This can be an excellent entry system save/ restore option and a good alternative to DAT72, DAT160, 8mm, and VXA-2 and VXA-320 tapes. CCIN: 63B8-005.

#EUA4 attaches to a Power server via a USB cable which carries data and control information. It is not powered by the USB port on the Power System or Power System USB adapter, but has a separate electrical line cord. Physically the #EUA4 docking station is a standalone enclosure about 2.0 x 7.0 x 4.25 inches in size which can sit on a shelf or on top of equipment in a rack.

#EUA4 is a follow on product to the #EU04 RDX docking station. #EUA4 has identical performance and identical application function to:

- Previously announced #EU04 and #1104 USB external docking stations
- Top mount #EUA3 USB docking station used in the Power S814 tower configuration
- #EU03 USB internal docking stations used in Power 720 system units
- #EU03 USB docking station available in the rack mounted IBM 7226-1U3 Multimedia Drawer

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

- Attributes provided:
  - USB 3.0 RDX External Docking Station
  - USB 3.0 Cable (2.7 meter or 8.8 foot)
  - Four line cords (1.85 meter or 6 foot) with type A, G, F or I plug (see http://www.iec.ch/worldplugs for type definition and country- specific usage)
  - One power jumper cord as an alternative to using one of the four power line cords above. This would draw power from a PDU in a rack.
  - Power Adapter using single phase 110-250V 50-60Hz power source
- Attributes required:
  - One USB port on server or server's USB adapter
  - At least one Removable Disk Drive Cartridge such as #EU01 or #1107
- Minimum required: 0
- Maximum allowed: 6 (Initial order maximum: 6)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EUA5) Standalone USB DVD drive w/cable

The Standalone USB DVD drive (FC EUA5) is an optional, standalone external USB-DVD device. It requires high current at 5V and must use the front USB 3.0 port on the 9008-22L, 9009-22A, 9009-41A, 9009-42A, 9223-22H, and 9223-42H systems.

- Custom card identification number (CCIN): 6331 model 005
- Media: Reads CD-ROM, CD-R, CD-RW, DVD-R, DVD+R, DVD-RW, DVD+RW, DVD-ROM, and DVD-RAM discs; Writes 4.7 GB DVD-RAM discs; CD read speed at 24X; DVD-ROM read speed at 8X; DVD-RAM at 5X; DVD-RAM has a write speed of 5X; The buffer size is 0.75 MB and cannot be disabled.
- Interface: USB
- Connector: USB 2.0
- Loading tray: supports 12 cm and 8 cm discs floor) Note: A USB extension cable is included (P/N 32N1311). The USB extension cable is to be used when there are no safe, flat spots available in the rack. This cable allows the drive to reach the floor. Alternate or additional extension cables are not supported as the total USB cable length can be no longer than 3 meters.
- Form factor: standalone USB DVD drive
- DVD video: not supported
- Attributes provided: USB DVD drive
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required:
  - AIX - supported
  - IBM i - supported
  - Linux - supported

  Refer to Software Requirements for specific O/S levels supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No
  Note: assignment to the VIOS supported

(#EVSN) - Enable Virtual Serial Number

This feature provides the ability to order a set of Virtual Serial Numbers.

The order is placed and fulfilled on the Entitled System Support (ESS) web site. A customer can go to ESS and request their set of Virtual Serial Numbers on the machine. At the end of the ordering and ESS retrieval steps, the customer will have an array of Virtual Serial Numbers available to start using, example: assigning to partitions, on their machine.

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 999 (Initial order maximum: 999)
- OS level required: IBM i supported
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#EXA2) 2 YEAR, ADVANCED EXPERT CARE

This feature indicates the IBM Power Expert Care Advanced 2 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)

- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EXA3) 3 YEAR, ADVANCED EXPERT CARE

This feature indicates IBM Power Expert Care Advanced 3 years

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: Yes
- Return parts MES: No

(#EXA4) 4 YEAR, ADVANCED EXPERT CARE

This feature indicates the IBM Power Expert Care Advanced 4 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EXA5) 5 YEAR, ADVANCED EXPERT CARE

This feature indicates IBM Power Expert Care Advanced 5 years

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: Yes
- Return parts MES: No

(#EXP1) 1 YEAR, PREMIUM EXPERT CARE

This feature indicates IBM Power Expert Care Premium 1 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: Yes
- Return parts MES: No

(#EXP2) 2 YEAR, PREMIUM EXPERT CARE

This feature indicates the IBM Power Expert Care Premium 2 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EXP3) 3 YEAR, PREMIUM EXPERT CARE

This feature indicates IBM Power Expert Care Premium 3 years

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: Yes
- Return parts MES: No

(#EXP4) 4 YEAR, PREMIUM EXPERT CARE

This feature indicates the IBM Power Expert Care Premium 4 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EXP5) 5 YEAR, PREMIUM EXPERT CARE

This feature indicates IBM Power Expert Care Premium 5 years

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: Yes
- Return parts MES: No

(#EXA2) 2 YEAR, ADVANCED EXPERT CARE

This feature indicates the IBM Power Expert Care Advanced 2 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EXA4) 4 YEAR, ADVANCED EXPERT CARE

This feature indicates the IBM Power Expert Care Advanced 4 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EXP2) 2 YEAR, PREMIUM EXPERT CARE

This feature indicates the IBM Power Expert Care Premium 2 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#EXP4) 4 YEAR, PREMIUM EXPERT CARE

This feature indicates the IBM Power Expert Care Premium 4 year

- Minimum required: 0
- Maximum allowed: 1 (Initial order maximum: 1)
- OS level required: None
- Initial Order/MES/Both/Supported: Initial
- CSU: N/A
- Return parts MES:

(#SVBP) BP Post-Sale Services

1 Day

BP Post-Sale Services: 1 Day

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 250 (Initial order maximum:)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#SVCS) IBM Systems Lab Services Post-Sale Services

1 Day

IBM Systems Lab Services Post-Sale Services: 1 Day

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 250 (Initial order maximum:)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#SVNN) Other IBM Post-Sale Services

1 Day

Other IBM Post-Sale Services: 1 Day

- Attributes provided: None
- Attributes required: None
- Minimum required: 0
- Maximum allowed: 250 (Initial order maximum:)
- OS level required: None
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: No

(#SVPC) 5000 Power to Cloud Reward points

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

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

For additional details, visit ibm.biz/PowertoCloud

This feature is ONLY for new initial orders or MES upgrades from POWER8 to POWER9. It DOES NOT allow loose part MES on POWER9

- Attributes provided: 5,000 Power to Cloud Reward points
- Minimum required: 0
- Maximum allowed: 2 (Initial order maximum: 2)
- OS level required:
- Initial Order/MES/Both/Supported: Both
- CSU: Yes
- Return parts MES: Does not apply
  Note: Feature SVPC not orderable in China

Feature exchanges

Not applicable.

↑ Back to top

Accessories

None.

Customer replacement parts

None.

Back to top

Supplies

None.

Supplemental media

None.

Back to top

Trademarks

(R), (TM), * Trademark or registered trademark of International Business Machines Corporation.

** Company, product, or service name may be a trademark or service mark 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 Sales Manual at any time without notice. This Sales Manual is provided for your information only. Additional terms of use are located at

Terms of use