

z196 is designed and optimized for growth and large-scale consolidation. Z196 provides:

- Up to 80 customer processor units (PUs)
 - 5.2 GHz high frequency z196 quad core processor chip.
 - Up to 768 GB of memory per book
 - Up to 15 PUs available for subcapacity use
- Maximum available real memory:**
- Up to 3.0 TB of maximum available memory available on z196 (depending on the model)
 - Up to 1.0 TB of maximum available memory supported on z196 in any LPAR.

Up to 3 TB of Redundant array of independent memory (RAIM)

- RAIM technology provides protection at the dynamic random access memory (DRAM), dual inline memory module (DIMM), and memory channel level.

IBM zEnterprise BladeCenter Extension (zBX) support

- zBX is a separate machine, machine type 2458 Model 002, attached to z196 and can be viewed as a logical extension of z196.
- With this support, heterogeneous applications distributed across multiple environments can be configured and processed in a single zEnterprise environment.
- With zBX, you can process CPU intensive DB2® queries, found in Business Intelligence and Data Warehousing applications, using the IBM Smart Analytics Optimizer for DB2 for z/OS, V1.1. With zBX, you can also run multiple instances of operating systems or virtual servers simultaneously on POWER7® blades optimizing workload performance and capacity.

IBM zEnterprise Unified Resource Manager (Unified Resource Manager)

Unified Resource Manager is part of the HMC that manages an ensemble. This includes providing energy monitoring and management, goal-oriented policy management, increased security, virtual networking, virtual server lifecycle management, and data management for the physical and logical resources of an ensemble.

- An ensemble is a collection of one to eight zEnterprise CPCs, including any optionally attached zBX, that are managed as a single logical virtualized system by the Unified Resource Manager.

Reduction in the preplanning requirements by:

- Providing a fixed HSA (16 GB)
- Reducing the number of Power-on-Resets
- Allowing dynamic add/remove of a new LPAR to a new or existing LCSS.

Coupling using InfiniBand®

- A zEnterprise to zEnterprise, zEnterprise to System z10, or System z10 to System z10 connection is provided by an InfiniBand Double Data Rate (12x IB-DDR) fiber optical link with a link data rate of 6 Gigabytes per second (Gbps) and maximum link distance of 150 meters (492 feet), or by a 1x IB-SDR fiber optical link or 1x IB-DDR fiber optical link with a maximum unrepeated distance of 10 kilometers (6.2 miles) and a maximum repeated distance of 100 kilometers (62 miles). 1x IB-SDR has a link data rate of 2.5 Gbps (Gigabits). 1x IB-DDR has a link data rate of 5 Gbps (Gigabits).

- A zEnterprise to System z9® or System z10 to System z9 connection is provided by an InfiniBand Single Data Rate (12x IB-SDR) fiber optical link with a link data rate of 3 Gigabytes per second (Gbps) and maximum link distance of 150 meters (492 feet).

InterSystem Channel-3 (ISC-3) The ISC-3 link is a member of the family of coupling link options. ISC-3 has a link data rate of 2 Gbps.

- It is used by coupled servers to pass information back and forth over high speed links in a Parallel Sysplex environment for unrepeated distances up to 10 km (6 miles) and repeated distances up to 100 km (62 miles).

IBM System z Integrated Information Processor (zIIP)

- zIIP is a specialty engine designed to help free-up general computing capacity and lower overall total cost of computing for select data and transaction processing workloads. Using a zIIP can help free capacity on the general-purpose processor.

IBM System z Application Assist Processor (zAAP)

- zAAP is a specialized processor unit that provides a strategic Java execution environment, which enables clients to integrate and run new Java-based web applications alongside core z/OS business applications and backend database systems.

Integrated Facility for Linux (IFL)

- An IFL is a specialty engine that provides additional processing capacity exclusively for Linux on System z workloads.

Internal Coupling Facility (ICF)

- An ICF is a specialty engine that provides additional processing capability exclusively for the execution of the Coupling Facility Control Code (CFCC) in a coupling facility partition.

HCA (Host Channel Adapter) fanout cards (HCA2-C (copper), HCA2-O (optical), and HCA2-O LR (optical))

- HCA2-C (copper) cards have InfiniBand connections used for internal I/O on zEnterprise. HCA2-O and HCA2-O LR (optical) cards have InfiniBand connections used for coupling on a zEnterprise. HCA2-O supports 12x InfiniBand at 3 Gbps or 6 Gbps. HCA2-O LR supports 1x InfiniBand at 2.5 Gbps or 5 Gbps.

Server Time Protocol function provides:

- The only time synchronization for zEnterprise using the FSP/STP card.
- Multisite sysplex distance to 100 km.
- Coexistence of non-zEnterprise servers and coupling facilities (CFs) synchronized in an ETR network with servers and CFs that are synchronized with Coordinated Server Time (CST)
- Concurrent migration from an ETR network
- Exchange of timing signals over ISC-3 links and InfiniBand (IFB) links NTP client support. The NTP client attaches to an NTP server provides time accuracy across heterogeneous platforms in an enterprise
- Enhanced accuracy to an external time source using pulse per second (PPS) output from NTP server
- Use of the HMC as an NTP server configured for use as the external time source
- Continuous availability of NTP servers used as an external time source
- Enhanced STP recover when the Internal Battery Feature is in use
- Ability to save the STP configuration and time information across Power on
- Resets (POR) or power outages for a single or dual server STP-only CTN
- Automation of STP CTN reconfiguration using the System z application programming interface (API)
- Ability to notify z/OS when events related to accessing an external time source occur.

Internet Protocol Version 6 (IPv6) support

- IPv6 is available for the HMC and SE customer network, the TKE network connection to operating system images, OSA-Express3, OSA-Express2, and HiperSockets™. IPv6 is the protocol designed by the Internet Engineering Task Force (IETF) to replace Internet Protocol Version 4.

Power estimating and monitoring functions:

- Power Estimator tool on Resource Link
- Monitoring of power consumption and thermal loading using the **Activity** task and the **Monitors Dashboard** task on the HMC
- Support for IBM Systems Director Active Energy Manager™ for x86, IBM Systems Director Active Energy Manager for POWER®, and IBM Systems Director Active Energy Manager for Linux on System z, which can monitor power and thermal data for zEnterprise, as well as other systems.

Future enhancements available for z196 models may be dependent on the EC level of the Central Processor Complex (CPC), Microcode Level (MCL) and/or HMC.



See zEnterprise 196 Installation Manual for Physical Planning, available on Resource Link at <http://www.ibm.com/servers/resourcelink>

Historical view of power, temperature, and utilization data of your system

- Using the **Environment Efficiency Statistics** task on the HMC and Support Element, you can display a historical view of system power consumption, system temperature, blade CPU utilization, and CP utilization data. This data will assist you in monitoring the performance of your system.

Energy consumption reduction

- You can reduce the energy consumption of your system by enabling power saving mode or setting a peak power consumption limit. To enable power saving mode, use the **Set Power Saving, Customize Schedule Operations, or Customize/Delete Activation Profiles** HMC and Support Element tasks, the Active Energy Manager (AEM), or SNMP and CIM APIs. To limit the peak power consumption, use the **Set Power Cap** task.

Capacity on Demand functions, which include:

- Ability to perform a permanent LICCC upgrade while temporary resources are active
- Ability to install and activate multiple temporary records at any given time
- Ability to activate partial resources on a single temporary record
- Disaster recovery solutions:
 - Capacity for Planned Events (CPE) - Short range - 3 days
 - Capacity Backup (CBU) - Long range - 90 days
- Capacity provisioning, which provides a means of managing your processing capacity based on business needs
- Ability to prepay for On/Off CoD upgrades
- Ability to set spending limits when ordering an On/Off record
- Ability to order permanent unassigned engines
- Ability to order an administrative On/Off test record, which allows you to order, download, activate, and deactivate On/Off upgrades without actually setting real capacity or incurring costs
- Automatic renewal of On/Off CoD records
- Automatic installation of up to four CPE and CBU records on an initial z196 order
- 45 available subcapacity settings.

HiperSockets provides the ability to diagnose network problems using the Network Traffic Analyzer (HS NTA). A single logical partition can connect up to 32 HiperSockets.

HiperDispatch helps provide increased scalability and performance of higher n-way and multi-book z196 systems by improving the way workload is dispatched across the server. HiperDispatch accomplishes this improvement by recognizing the physical processor where the work was started and then dispatching subsequent work to the same physical processor. This intelligent dispatching helps reduce the movement of cache and data, and improves CPU time and performance. Support to dynamically optimize the CPU-to-book allocation of physical processor (PUs). Available with zEnterprise PR/SM™, System z10 PR/SM, and z/OS functions.

Large page support (1 MB pages) provides performance improvement for a select set of applications, primarily long running memory access intensive applications.

Reduced impact of planned and unplanned server outages through:

- Enhanced book availability
 - Program directed re-IPL
 - Dynamic oscillator switchover
 - Redundant I/O interconnect
 - Enhanced driver maintenance
 - System-initiated CHPID reconfiguration
 - Concurrent HCA fanout card hot-plug and rebalance
- Enhanced driver maintenance** allows Licensed Internal Code (LIC) updates to be performed in support of new features and functions. When properly configured, z196 is designed to support activating a selected new LIC level concurrently. Certain LIC updates are not supported by this function.

Dynamic oscillator switchover

z196 has two oscillator cards, a primary and a backup. If a primary card failure occurs, the backup card is designed to detect the failure, switch over, and provide the clock signal to the server transparently.

Enhanced book availability allows a single book, in a multibook server, to be concurrently removed from the server and reinstalled during an upgrade or repair action. Enhanced book availability is an extension of the support for Concurrent Book Add (CBA).

Flexible memory provides the additional resources to maintain a constant level of memory.

Redundant I/O interconnect helps maintain critical connections to devices. z196 allows a single book, in a multibook server, to be concurrently removed and reinstalled during an upgrade or repair, continuing to provide connectivity to the server I/O resources using a second path from a different book.

Up to 60 logical partitions (LPARs)

Server consolidation The expanded capacity and enhancements to the I/O infrastructure facilitates the consolidation of multiple servers into one z196 with increased memory and LPARs, which might allow you to reduce the number of servers while hosting additional applications. z196 provides the ability to define up to four logical channel subsystems (LCSS). Each LCSS can support up to 256 CHPID definitions and 15 LPARs (up to a maximum of 60 LPARs per system).

Water-cooling option

- A water-cooling option is available to cool the server using customer supplied chilled water and a special water circulation unit within the frame.

Top exit I/O cabling

z196 provides the ability to route all I/O cables, ESCON, FICON, OSA-Express, 12x InfiniBand, 1x InfiniBand, and ISC-3, as well as 1000BASE-T Ethernet copper cabling from I/O cages or I/O drawers out through the top of the frame.

Frame bolt-down kits

- Two optional bolt-down kits (one for refrigerated/air-cooled models and one for water-cooled models) are available to help secure the frames and its contents from damage when exposed to vibrations and shocks. The kits supply parts to cover raised floor heights from 9-13 inches, 12-22 inches, and 12-36 inches.

High Voltage DC universal input option

- Ability to operate z196 using high voltage DC power (380-570 volts) in addition to AC power. The direct high voltage DC design improves data center energy efficiency by removing the need for any conversion.

ESCON (16 ports) supporting 240 channels

- **FICON Express8 and FICON Express4 Note:** FICON Express4 features can only be carried forward.
- FICON Express8 features support a link data rate of 2, 4, or 8 Gbps autonegotiating with the attached device. Features include:
 - FICON Express8 10KM LX (4 channels per feature)
 - FICON Express8 SX (4 channels per feature)
- FICON Express4 features support a link data rate of 1, 2, or 4 Gbps autonegotiating with the attached device. Features include:
 - FICON Express4 10KM LX (4 channels per feature)
 - FICON Express4 4KM LX (4 channels per feature)
 - FICON Express4 SX (4 channels per feature)

Enhancements:

- High Performance FICON for System z (zHPF) for FICON Express8 and FICON Express4 features (CHPID type FC)
- Extension to zHPF multitrack operations removing the 64 kB data transfer limit
- Assigning World Wide Port Names (WWPNs) to physical Fibre Channel Protocol (FCP) ports using the WWPN prediction tool


continued

OSA-Express3 and OSA-Express2

Note: OSA-Express2 features can only be carried forward.

- OSA-Express3 features:
 - OSA-Express3 GbE LX (4 ports per feature)
 - OSA-Express3 GbE SX (4 ports per feature)
 - OSA-Express3 1000BASE-T Ethernet (4 ports per feature)
 - OSA-Express3 10 GbE LR (2 ports per feature)
 - OSA-Express3 10 GbE SR (2 ports per feature)
- OSA-Express2 features:
 - OSA-Express2 GbE LX (2 ports per feature)
 - OSA-Express2 GbE SX (2 ports per feature)
 - OSA-Express2 1000BASE-T Ethernet (2 ports per feature)
- Enhancements:
 - OSA-Express3 CHPID type OSM provides connectivity to the intranode management network (INMN)
 - OSA-Express3 CHPID type OSX provides connectivity and access control to the intraensemble data network (IEDN)
 - Inbound workload queuing (IWQ) for OSA-Express3

Resource Link is a key component in getting your z196 system up and running and maintained. Key areas of improvement include the customized planning aids, CHPID Mapping Tool, Customer Initiated Upgrade (CIU), power estimation tool, and additional education courses.



IBM's Large Systems Performance Reference (LSPR) method provides comprehensive z/Arch. processor capacity data for different configurations of Central Processor Units across a wide variety of system control program & workload environments.

Cryptographic options:

- Configurable Crypto Express3 feature.
- CP Assist for Cryptographic Function (CPACF), which delivers cryptographic support on every PU with data encryption/decryption. CPACF also provides a high performance secure key function that ensures the privacy of key material used for encryption operations.
 - CPACF support includes AES for 128-, 192-, and 256-bit keys; SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 for message digest; PRNG, DES, and TDES.
- Using the Support Element, you can enable or disable the encrypt DEA key and encrypt AES key functions of the CPACF.
 - Elliptic Curve Cryptography (ECC) and RSA public-key cryptography support
 - User Defined Extension (UDX) support.
 - Remote loading of ATMs and POS keys.
 - Dynamically add, move, or delete a Crypto Express3 feature to or from an LPAR.
 - Cryptographic migration wizard on TKE for migrating configuration data from one Cryptographic coprocessor to another Cryptographic coprocessor.
 - The tamper-resistant hardware security module, which is contained within the Crypto Express3 is designed to meet the FIPS 140-2 Level 4 security requirements for hardware security modules.

Fiber Quick Connect (FQC), an optional feature, is a fiber harness integrated in the z196 frame for a "quick" connect to ESCON and FICON LX channels.

SNMP Client Libraries 3.0 support

CFCC Level 17 support

TKE 7.0 Licensed Internal Code (LIC) support

z/VM@-mode partition (LPAR) support to contain processor types (CPs, IFLs, zIIPs, zAAPs, and ICFs)

Plan ahead memory is an optional feature, allows you to preplan to future memory upgrades. The memory upgrades can be made nondisruptively and also concurrently.

Worldwide Port Name (WWPN) prediction tool

- The WWPN prediction tool assists you in preplanning and setting up your Storage Area Networks (SANs) environment. The WWPN prediction tool can calculate WWPNs for virtual and physical ports before the installation of z196.
 - Therefore, you can be up and running much faster after the server is installed.
 - This tool applies to all FICON channels defined as CHPID type FCP (for communication with SCSI devices). The WWPN prediction tool is located on Resource Link.

Support to control user access to the HMC using a pattern name that defines:

- Search criteria used to identify specific user IDs
- LDAP server used for authentication
- HMC user ID template used to identify logon permissions for the user IDs using this template
- List of HMCs that can be accessed.

Enhanced security using digital signatures

- Digitally Signed Firmware (Licensed Internal Code) support provided by the HMC and the SE. This support provides the following benefits:
 - Ensures that no malware can be installed on System z products during firmware updates (such as, transmission of MCL files, delivery of code loads, and restoration of critical data)
 - Designed to comply to FIPS (Federal Information Processing Standard) 140-2 Level 1 for Cryptographic LIC (Licensed Internal Code) changes.

Auditability function

- HMC/SE tasks are available to generate, view, save, and offload audit reports, to set up a schedule for generating, saving, and offloading audit information, to receive email notifications for select security log events, and to remove the predefined password rules to prevent them from being mistakenly used.
 - You can also manually offload or set up a schedule to automatically offload HMC and Support Element log files, which can help satisfy audit requirements.

This section lists the features/functions that are not supported on z196.

- OSA-Express is not supported on z196.
- OSA-Express2 10 GbE LR is not supported on z196.
- Crypto Express 2 is not supported on z196.
- FICON Express and FICON Express2 are not supported on z196.
- The ETR feature is not supported on z196.
- ICB-4 links are not supported on z196.
- ICB-3 links are not supported on z196.
- ICB-2 links are not supported on z196.
- ISC-3 links in compatibility mode are not supported on z196 (CHPID types CFS and CFR).
- PCIXCC and PCICA are not supported on z196.

z196 (machine type 2817) is offered in five models.

- The model naming is representative of the maximum number of customer configurable processor units (PUs) in the system. PUs are delivered in single engine increments orderable by feature code.
- The model number also reflects the number of books installed.

With the expanded capacity of z196 and enhancements to the I/O infrastructure, IBM continues to facilitate the consolidation of multiple servers into one z196 with A substantial increase in:

- Available memory
- Advanced virtualization
- LPARs
- Speed using InfiniBand
- Available processors in a single footprint

Model number	Maximum number of books allowed	Refrigerated/Air-cooled feature code	Water-cooled feature code
M15	One book	FC 1125	FC 1130
M32	Two books	FC 1126	FC 1131
M49	Three books	FC 1127	FC 1132
M66	Four books	FC 1128	FC 1133
M80	Four books	FC 1129	FC 1134

- The following table lists the five z196 models and some of their characteristics, such as range of PUs allowed, the memory range of each model, and number of InfiniBand connections, and the range of I/O cages and I/O drawers that can be installed. **NOTE:** These values are the same for both the refrigeration/air-cooled and the water-cooled models.

Models	Processor Units (PUs)	Memory ¹	InfiniBand connections ²	I/O cages ³	I/O drawers
M15	1 to 15	32 GB to 704 GB	16	0 to 2	0 to 4
M32	1 to 32	32 GB to 1520 GB	32	0 to 2	0 to 6
M49	1 to 49	32 GB to 2288 GB	40	0 to 2	0 to 6
M66	1 to 66	32 GB to 3056 GB	48	0 to 2	0 to 6
M80	1 to 80	32 GB to 3056 GB	48	0 to 2	0 to 6

References:

- A portion of this memory (16 GB) is delivered and reserved for HSA.
- With the addition of the third and fourth books, air flow concerns require that you reduce the number of fanout cards plugged to increase cooling.
- With RPQ 8P2506, up to three I/O cages can be ordered.

- The system model number does not reflect the number of processors that are enabled for use. It reflects only the maximum number of customer-used processors that can be enabled when purchased.
 - The CP features offered have varying levels of capacity. The capacity setting is based on the quantity and type of CP feature.
 - It is identified by a **model capacity indicator**. The model capacity indicator identifies the number of active CPs rather than the total physical PUs purchased and identifies the type of capacity.
 - The model capacity indicators are identified as 7xx, 6xx, 5xx, and 4xx, where xx is the number of active CP features.
 - 7xx is a full capacity identifier. 6xx, 5xx, and 4xx are subcapacity identifiers.
- NOTE:** Subcapacity CP features provide reduced capacity relative to the full capacity CP feature.
 - While you can only have up to 15 subcapacity CPs, you are not limited to one book.
 - Example, model capacity indicator "700" indicates no active CPs at full capacity; Model capacity indicator "510" indicates 10 active CP5 type PUs.

	PU Type (Capacity Identifier)	Allowable Quantity (xx =) ¹		
Full capacity	CP7 (7xx)	00 - 15		
		00 - 32		
		00 - 49		
		00 - 66		
		00 - 80		
Subcapacity	CP6 (6xx)	1 - 15 ²		
		Subcapacity	CP5 (5xx)	1 - 15 ²
				Subcapacity

References:

- Depends on the 2817 machine model.
- For all 2817 models.

z/Architecture on z196 has:

- High-frequency z196 quad core processor chip (5.2 GHz operation in system)
- Out-of-order execution of instructions
- Hardware accelerators on the chip for data compression, cryptographic functions, and decimal floating point
- Integrated SMP communications
- Instructions added to the z196 chip to improve compiled code efficiency
- Enablement for software/hardware cache optimization
- z196 support for 1 MB page frames
- Full hardware support for Hardware Decimal Floating-point Unit (HDFU)
- 64-bit general registers
- 64-bit integer instructions. Most ESA/390 architecture instructions with 32-bit operands have new 64-bit and 32- to 64-bit analogs
- 64-bit addressing is supported for both operands and instructions for both real addressing and virtual addressing
- 64-bit address generation. z/Architecture provides 64-bit virtual addressing in an address space, and 64-bit real addressing.
- 64-bit control registers. z/Architecture control registers can specify regions and segment, or can force virtual addresses to be treated as real addresses
- The prefix area is expanded from 4 K to 8 K bytes
- Quad-word storage consistency
- The 64-bit I/O architecture allows CCW indirect data addressing to designate data addresses above 2 GB for both format-0 and format-1 CCWs.
- The 64-bit SIE architecture allows a z/Architecture server to support both ESA/390 (31-bit) and z/Architecture (64-bit) guests and Zone Relocation is expanded to 64-bit for LPAR and z/VM
- 64-bit operands and general registers are used for all cryptographic instructions
- The implementation of 64-bit z/Architecture can help reduce problems associated with lack of addressable memory by making the addressing capability unlimited (16 Exabytes)