



IBM zEnterprise 114 - Freedom by design

Table of contents

| | |
|---|--|
| 2 Overview | 39 Product number |
| 5 Key prerequisites | 49 Publications |
| 5 Planned availability date | 50 Technical information |
| 6 Description | 63 Pricing |
| 36 Product positioning | 64 Order now |
| 37 Statement of general direction | 65 Corrections |

At a glance

At the core of the IBM® zEnterprise™ System for both midsize and small enterprises is the IBM zEnterprise 114 (z114) - a true "system of systems" bringing the unique value of hybrid computing to a much broader set of businesses. The z114 is part of a workload-optimized multiarchitecture computer system capable of integrating and centrally managing multitier applications running across z/OS®, z/VM®, z/VSE™, z/TPF, Linux® on System z®, AIX®, and Linux on IBM System x® operating environments. This integration extends the gold standards of the mainframe's superior manageability, governance, and service quality to workloads connected to IBM System z using the IBM zEnterprise BladeCenter® Extension (zBX) Model 002 and managed by the IBM zEnterprise Unified Resource Manager.

The z114 continues the heritage in mainframe qualities of service with increased flexibility and performance in a lower-cost package. As an entry-level enterprise server with extreme granularity and extensive growth options, the z114 is designed to deliver significant improvements in packaging, performance, and total system scalability over prior generations. The z114 offers unprecedented scalability for small and medium businesses, providing more than a 15% improvement in performance as compared to its predecessor, the IBM System z10® Business Class.

The IBM zEnterprise 114 is designed to deliver:

- Innovative hybrid computing for new workloads and application deployment
 - A true "data center in a box" to host and centrally manage multiplatform workloads with data hosted on System z
 - A fit-for-purpose system optimized for today's leading application and database workloads
 - An expanded application portfolio that includes not just the traditional environments, but also AIX, Linux on System x, and Linux on System z workloads
 - Support for the IBM Smart Analytics Optimizer for DB2® for z/OS, V1.1, the IBM WebSphere® DataPower® Integration Appliance XI50 for zEnterprise (DataPower XI50z), and select POWER7® and IBM System x blades
- Server improvements
 - One to two drawers (system boards)
 - Seven processing units (PUs) per drawer
- Price/performance improvements across the stack
 - Lowest cost of entry ever
 - Extreme granularity with 130 available capacity settings, providing the freedom to choose the right capacity setting for your needs with the flexibility to scale on demand as workload demands increase

- Significant reduction in maintenance, hardware, specialty engine, and memory pricing, making new workload deployment even more affordable
- Improved hardware and software performance
 - Quad-core 3.8 GHz processor chips, and new out of order execution sequence designed to deliver world-class per-thread performance
 - Up to 14 cores with 10 that are user-configurable and up to 2 dedicated spares
 - Improvements of up to 18% in uniprocessor performance and 12% in total system capacity for z/OS , z/VM , and Linux workloads on System z compared to the z10™ BC
 - Over 100 new instructions and compiler-related enhancements that may help drive performance improvements of up to 25%, based on measurements and projections, for CPU-intensive workloads when accompanied by multiple C/C++ compiler-level improvements going from XL C/C++ V1R9 to V1R12
 - New PCIe I/O subsystem infrastructure that delivers faster and higher I/O bandwidth and increased I/O connectivity for storage networking and coupling
 - Up to 256 GB of available real random array of independent memory (RAIM)-protected fault-tolerant memory per server (System z exclusive) for improved availability and growing application needs, including an 8 GB fixed hardware system area (HSA)
 - Crypto Express3 cryptographic enhancements
 - Up to 128 coupling channel path identifiers (CHPIDs) - twice the number on the z10 BC - and up to 72 physical external coupling links (12x InfiniBand, 1x InfiniBand, ISC-3), an increase of almost 30% as compared to the z10 BC
 - Up to 240 ESCON® channels
- Support for future data center modernization, design, and efficiencies
 - Single frame, air cooled system
 - More performance to help you run more energy efficiently in the same footprint
 - Support for either top- or bottom-exit I/O and power cabling and raised-floor and non-raised-floor options
 - Optional high-voltage DC power, which can save you 1% - 3% on average, on your power bills
 - Humidity and altimeter smart sensors, which provide power and thermal optimization management

For ordering, contact your IBM representative, an IBM Business Partner, or IBM Americas Call Centers at 800-IBM-CALL (Reference: YE001).

Overview

IBM zEnterprise 114 (z114): Delivering the unique capabilities and value of hybrid computing to organizations of all sizes

With the announcement of the IBM zEnterprise System, IBM took smarter computing and optimized systems to a new level. For the first time it was possible to deploy an integrated hardware platform that brings mainframe and distributed technologies together, a system that can start to replace individual islands of computing and that can work to reduce complexity, improve security, and bring applications closer to the data they need. zEnterprise represents the next generation of smarter computing - hybrid computing where hundreds of workloads can run across multiple hardware platforms, all as part of a single integrated system.

Today IBM is introducing the IBM zEnterprise 114 (z114), the newest member of the zEnterprise family, bringing the unique value of hybrid computing to a much broader set of businesses. The z114 offers a smaller mainframe footprint and lower entry cost than the zEnterprise 196 (z196) and continues to deliver world-class secure data serving and transaction processing capabilities as well as the ability to help collapse and simplify infrastructures and efficiently host private enterprise clouds. The z114 delivers the scalability, flexibility, and performance you need at a lower capacity and attractive entry price point for your business.

Its hybrid capabilities are designed to address the complexity and inefficiency in today's multiarchitecture data centers. As a true "system of systems," the z114 extends the strengths and capabilities of the mainframe such as virtualization, security, fault tolerance, efficiency, and dynamic resource allocation to integrate and manage multitier applications running across z/OS , Linux on System z , VSE, z/VM , AIX , and Linux on System x operating environments, fundamentally changing the way data centers can be managed.

A unified system for integrated workloads: The zEnterprise System includes a central processor complex (CPC), either the z196 or the z114, the zEnterprise BladeCenter Extension (zBX) with its integrated optimizers and/or select IBM blades, and the zEnterprise Unified Resource Manager (Unified Resource Manager).

With the zBX and the zEnterprise Unified Resource Manager, you can combine System z , System p® , and System x server technologies into a single unified system, integrating workloads with affinity to mainframe applications and data, and manage it all with the same tools, techniques, and resources for consistent, automated, and reliable service delivery. This capability provides the unique ability to extend the strategic role of the mainframe and an opportunity to simplify and reduce the range of skills necessary for managing the IT infrastructure.

Now businesses can have a single, integrated view of the entire workload or business process. This revolutionary yet simplified multiarchitecture design helps drive operational efficiencies in deployment, integration, and management. It is designed to enhance the attractiveness of hosting applications on the mainframe by providing the freedom and ease of selecting the right architecture and operating system for application deployment within a common management infrastructure for both mainframe and distributed system resources.

Whether you are looking to expand delivery of new applications or cloud services, grow business transactions without growing IT costs, consolidate your infrastructure for reduced complexity, or address the inefficiencies of today's multitier workloads the z114 can help you handle your most complex and demanding business requirements, now and in the future.

Continued qualities of service: The z114 continues the heritage in mainframe qualities of service with extreme granularity, structure, and significant improvements in packaging, performance, and total system scalability over prior generations.

More processing capacity and performance: System resources are powered by up to 14 microprocessors running at 3.8 GHz. The z114 provides up to an 18% improvement in uniprocessor performance and up to a 12% increase in total system capacity for z/OS , z/VM , z/VSE , and Linux on System z workloads as compared to its predecessor, the z10 BC.

Its superscalar microprocessor chip has a higher-frequency design that leverages IBM technology leadership with a new out-of-order execution sequence that delivers world-class per-thread performance. With over 100 new instructions and numerous compiler related enhancements, the z114 can deliver up to 25% performance improvement, based on measurements and projections, for CPU-intensive workloads when accompanied by multiple C/C++ compiler level improvements going from XL C/C++ V1R9 to V1R12.

Scalability and flexibility for growth: The z114 will be available in two models: a single central processing drawer model, the M05, and a two drawer model, the M10, which offers the additional flexibility for I/O and coupling expansion and/or increased specialty engine capability. With up to 10 configurable cores, the model naming is indicative of how many total processor units are available for user characterization. The cores can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processors (zIIPs), Internal Coupling Facilities (ICFs), or additional System Assist Processors (SAPs), or can be used as additional spares (M10 only).

With 130 available capacity settings and a granular cost structure offered across either model, you have the freedom to choose the right capacity setting for your needs with the flexibility to scale on demand as workload demands increase.

To help secure sensitive data and business transactions, the z114 is designed for Common Criteria Evaluation Assurance Level 5 (EAL5) certification for security of logical partitions. Support for the next generation of public key technologies is available with Elliptic Curve Cryptography (ECC), which is ideal for constrained environments such as mobile devices. The z114 also offers support for key ANSI and ISO standards for the banking and finance industry.

256 GB memory for application growth: The z114 will support up to 248 GB of real (usable) RAIM-protected memory, an industry exclusive currently available only on System z . Beyond the purchased memory, there is an additional 8 GB of memory for the hardware system area (HSA), which holds the I/O configuration data for the server.

High-speed access to your resources: High-speed connectivity is critical in achieving sufficient levels of transaction throughput and enabling resources inside and outside the server to maximize application performance. The host bus interface of the z114 is designed to help satisfy clustering, security, SAN, and LAN requirements. Security for Secure Sockets Layer (SSL) transactions and secure coprocessing is delivered with Crypto Express3. IBM now also offers a PCIe I/O drawer for FICON® and OSA-Express multimode and single mode fiber optic environments designed to deliver increased capacity, infrastructure bandwidth, and reliability.

Packaging: For ease of installation, the z114 is a single-frame, air-cooled system that now supports either top- or bottom-exit I/O and power, raised-floor and non-raised-floor options, and high-voltage DC power, providing increased flexibility to accommodate small data center installations and support for future data center design and efficiencies.

zEnterprise offers a revolutionary new IT operating environment that easily integrates and manages distributed technologies with mainframe resources, which can dramatically increase the flexibility of the IT infrastructure to reduce costs today and support future business growth. zEnterprise is smarter computing, with industry-leading management capabilities. It is perfectly positioned for cloud computing with heterogeneous resource pools, scalability, security, and policy-driven autonomies. It really is a virtual cloud or data center in a box!

The zBX can support the IBM Smart Analytics Optimizer for DB2 for z/OS , V1.1 (5697-AQT), the IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z), and select POWER7 and IBM System x blades.

Smart Analytics Optimizer - a specialized accelerator for queries: The IBM Smart Analytics Optimizer is a high-performance, integrated hardware and software specialized accelerator that delivers dramatically faster analytic response times to select queries, complementing traditional query processing. It offers business value at lower cost through the simplicity and rapid deployment characteristics of an appliance, while retaining the benefits of having the data managed and secured by DB2 for z/OS . This workload-optimized, appliance-like add-on is installed in the zBX, connects to DB2 transparently to the user and applications, and requires no changes to an existing application to take immediate advantage of the analyses.

The IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z) is a multifunctional appliance that can help provide multiple levels of XML optimization, streamline and secure valuable service-oriented architecture (SOA) applications, and provide drop-in integration for heterogeneous environments by enabling core enterprise service bus (ESB) functionality, including routing, bridging, transformation, and event handling. It can help to simplify, govern, and enhance the network security for XML and web services.

POWER7 blades (supporting AIX) and/or **System x blades** (supporting Linux) installed in the zBX can enable application integration with System z transaction processing, messaging, and data serving capabilities. The blades are managed

as part of a single logical virtualized environment by the IBM zEnterprise Unified Resource Manager.

The IBM zEnterprise Unified Resource Manager manages System z ensembles, collections of one or more zEnterprise System nodes in which each node is comprised of a zEnterprise (z114 or z196) and its optionally attached IBM zEnterprise BladeCenter Extension (zBX). An ensemble can consist of a single zEnterprise with no zBX attached, or up to eight CPCs where at least one of the CPCs has a zBX attached. The resources of a zEnterprise System ensemble are managed and virtualized as a single pool of resources, integrating system and workload management across the multisystem, multitier, multiarchitecture environment.

IBM Global Technology Services (GTS) and Systems Lab Services can help you assess and perform an application fit-for-purpose analysis, and design an IBM zEnterprise System that integrates your IT strategy and business priority. This includes developing the business case and high-level transition plan, and a roadmap for an adaptable and efficient infrastructure. GTS and Systems Lab Services can also enable you to build and run a smarter IBM zEnterprise System environment. With these services, you can migrate effectively and efficiently to an IBM zEnterprise System, create a more cost-effective and manageable computing environment with server, storage, and network optimization, integration, and implementation, and effectively run and manage the IBM zEnterprise System with maintenance and technical support services.

Key prerequisites

Refer to the [Hardware requirements](#) and [Software requirements](#) sections of this announcement.

Planned availability date

- September 9, 2011
 - IBM zEnterprise 114 Models M05 and M10 new builds
 - MES orders for the following features:
 - IBM System z9® BC 2096 and System z10 BC 2098 model conversions to zEnterprise 114 Models M05 and M10
 - Field installed features and conversions that are delivered solely through a modification to the machine's Licensed Internal Code (LIC)
 - Full Hardware Management Console (HMC) support for 9x driver
 - Full Trusted Key Encryption (TKE) support for 7.1 code
 - Entitlement features for POWER7 , IBM Smart Analytics Optimizer, and DataPower blades on a zBX
- September 26, 2011
 - Manage Firmware System x Blade (#0042)
 - Advanced Management Firmware System x Blade (#0046)
- October 21, 2011
 - Add a zBX as an MES to the installed z114
- December 31, 2011
 - MES orders for the following features:
 - M05 to M10 model conversions
 - M10 to z196 M15 model conversions
 - Loose piece MESs

Description

As a unified system for integrated workloads, the zEnterprise System includes:

- A central processor complex (CPC), either the zEnterprise 196 (z196) and / or the zEnterprise 114 (z114),
- The zEnterprise BladeCenter Extension (zBX) with its integrated optimizers and / or select IBM blades,
- The zEnterprise Unified Resource Manager (Unified Resource Manager).

With the zEnterprise BladeCenter Extension and the zEnterprise Unified Resource Manager, you can combine System z, UNIX®, and Intel® server technologies into a single unified system integrating workloads with affinity to mainframe applications and data and manage it all with the same tools, techniques, and resources for consistent, automated, and reliable service delivery. This capability provides the unique ability to extend the strategic role of the mainframe across the IT infrastructure and an opportunity to simplify and reduce the range of skills necessary for managing the infrastructure.

Now businesses can have a single, integrated view of the entire workload or business process. This revolutionary yet simplified multiarchitecture design helps drive operational efficiencies in deployment, integration, and management. It is designed to enhance the attractiveness of hosting applications on the mainframe by providing the freedom and ease of selecting the right architecture and operating system for application deployment within a common management infrastructure for both mainframe and distributed system resources.

Continued qualities of service: The z114 continues the heritage in mainframe qualities of service with a more granular cost structure and significant improvements in packaging, performance, and total system scalability over prior generations.

More processing capacity and performance: System resources are powered by up to 14 microprocessors running at 3.8 GHz. The z114 provides increases of up to 20% in performance per core and up to 30% in total system capacity for z/OS, z/VM, and Linux on System z workloads compared to its predecessor, the z10 BC.

The new microprocessor chip has a higher-frequency design that leverages IBM leadership technology with more cache than other chips and a new execution sequence that delivers world-class per-thread performance. There are 100 new instructions that will help to deliver CPU-centric performance. For CPU-intensive workloads, additional gains of up to 30% can be achieved via multiple compiler-level improvements.

Scalability for granular cost: The z114 will be available in two models: a single central processing drawer model, the M05, and a two drawer model, the M10, which offers additional flexibility for I/O and coupling expansion and/or increased specialty engine capability. With up to 10 configurable cores, the model naming is indicative of how many total processor units are available for user characterization. The cores can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processors (zIIPs), Internal Coupling Facilities (ICFs), or additional System Assist Processors (SAPs), or can be used as additional spares (M10 only).

And with 130 available capacity settings and a granular cost structure offered across either model, you have the freedom to choose the right capacity setting for your needs with the flexibility to scale on demand as workload demands increase.

To help secure sensitive data and business transactions, the z114 is designed for Common Criteria Evaluation Assurance Level 5 (EAL5) certification for security of logical partitions. And support for the next generation of public key technologies is available with Elliptic Curve Cryptography (ECC), which is ideal for constrained environments such as mobile devices. The z114 also offers support for key ANSI and ISO standards for the banking and finance industry.

More memory for application growth: The z114 will support up to 248 GB of real (usable) RAIM-protected memory, an industry exclusive currently available only on System z . Beyond the purchased memory, there is an additional 8 GB of memory for the hardware system area (HSA) which holds the I/O configuration data for the server.

High-speed access to your resources: High-speed connectivity is critical in achieving sufficient levels of transaction throughput and enabling resources inside and outside the server to maximize application performance. The host bus interface of the z114 is designed to help satisfy clustering, security, SAN, and LAN requirements. Security for Secure Sockets Layer (SSL) transactions and secure coprocessing is delivered with Crypto Express3. IBM now also offers a PCIe I/O drawer for FICON and OSA-Express multimode and single mode fiber optic environments designed to deliver increased capacity, infrastructure bandwidth, and reliability.

Packaging: For ease of installation, the z114 is a single-frame, air-cooled system that now supports either top- or bottom-exit I/O and power, raised-floor and non-raised-floor options, and high-voltage DC power, providing increased flexibility to accommodate small data center installations and support for future data center design and efficiencies.

zEnterprise offers a revolutionary new IT operating environment that easily integrates and manages distributed technologies with mainframe resources, which can dramatically increase the flexibility of the IT infrastructure to reduce costs today and support future business growth. zEnterprise is smarter computing, with industry-leading management capabilities. It is perfectly positioned for cloud computing - heterogeneous resource pools, scalability, security, and policy-driven autonomies. It really is a virtual cloud or data center in a box!

The zBX can support IBM Smart Analytics Optimizer for DB2 for z/OS , V1.1 (5697-AQT), the IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z), and select POWER7 and IBM System x blades.

Smart Analytics Optimizer - a specialized accelerator for queries: The IBM Smart Analytics Optimizer is a high-performance, integrated hardware and software specialized accelerator that delivers dramatically faster analytic response times to select queries, complementing traditional query processing. It offers business value at lower cost through the simplicity and rapid deployment characteristics of an appliance, while retaining the benefits of having the data managed and secured by DB2 for z/OS . This workload-optimized, appliance-like add-on is installed in the zBX, connects to DB2 transparently to the user and applications, and requires no changes to an existing application to take immediate advantage of the analyses.

The IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z) is the latest addition to the zEnterprise integrated infrastructure. DataPower XI50z is a multifunctional appliance that can help provide multiple levels of XML optimization, streamline and secure valuable service-oriented architecture (SOA) applications, and provide drop-in integration for heterogeneous environments by enabling core enterprise service bus (ESB) functionality, including routing, bridging, transformation, and event handling. It can help to simplify, govern, and enhance the network security for XML and web services.

The optional POWER7 blades (supporting AIX) and/or System x blades (supporting Linux) installed in the zBX can enable application integration with System z transaction processing, messaging, and data serving capabilities. The blades are managed as part of a single logical virtualized environment by the IBM zEnterprise Unified Resource Manager.

The IBM zEnterprise Unified Resource Manager manages System z ensembles - collections of one or more zEnterprise System nodes in which each node is comprised of a z114 and/or its optionally attached IBM zEnterprise BladeCenter Extension (zBX). An ensemble can consist of a single z114 with no zBX attached, or up to eight CPCs where at least one of the z114s has a zBX attached. The resources of a zEnterprise System ensemble are managed and virtualized as a single pool of

resources integrating system and workload management across the multisystem, multitier, multiarchitecture environment.

IBM Global Technology Services (GTS) and Systems Lab Services can help you assess and perform an application fit-for-purpose analysis, and design an IBM zEnterprise System that integrates your IT strategy and business priority. This includes developing the business case, the high-level transition plan, and a roadmap for an adaptable and efficient infrastructure. GTS and Systems Lab Services can also enable you to build and run a smarter IBM zEnterprise System environment. With these services, you can migrate effectively and efficiently to an IBM zEnterprise System, create a more cost-effective and manageable computing environment with server, storage, and network optimization, integration, and implementation, and effectively run and manage the IBM zEnterprise System with maintenance and technical support services.

With a total systems approach designed to deploy innovative technologies, IBM System z introduces the z114, supporting zEnterprise , and offering the highest levels of reliability, availability, scalability, clustering, and virtualization. The z114 just-in-time deployment of capacity allows improved flexibility and administration, and the ability to enable changes as they happen. The expanded scalability on the z114 facilitates growth and large-scale consolidation.

The IBM zEnterprise 114 is designed to provide:

- Server improvements
 - One to two drawers (system boards)
 - Seven PUs per drawer
 - Up to 30 logical partitions
 - Up to 248 GB per partition
- Memory
 - Up to 248 GB memory (Model M10)
 - 8 GB system minimum
 - Increments of 8 or 32 GB
 - 8 GB HSA, separately managed
 - RAIM standard
- I/O
 - Zero to two I/O drawers with concurrent add/repair capability
 - Zero to two PCIe I/O drawers with concurrent add/repair capability
- Coupling
 - HCA3-O 12x optical 2 port
 - HCA3-O LR 1x optical 4 port
 - Up to 48 ISC-3 links
 - Up to 16 12x PSIFB ports (Model M10)
 - Up to 32 1x PSIFB ports (Model M10)
- zBX support
 - zManager
 - BladeCenter H (1-8 chassis) Ethernet attachment
 - Support for IBM Smart Analytics Optimizer, DataPower IX50z, System p blades, and System x blades

The performance advantage

IBM's Large Systems Performance Reference (LSPR) method is designed to provide comprehensive z/Architecture® processor capacity ratios for different configurations of central processors (CPs) across a wide variety of system control programs and workload environments. For z114, the z/Architecture processor

capacity indicator is defined with a (A0x-Z0x) notation, where x is the number of installed CPs.

In addition to the general information provided for z/OS V1.11, the LSPR also contains performance relationships for z/VM and Linux operating environments.

The z114 (2818) processor is expected to provide an 18% improvement in performance as compared to a z10 BC (2098), based on workload and model. The largest z114 (2818-M10) is expected to provide 12% more capacity than the largest z10 BC (2098-E10).

The LSPR contains the Internal Throughput Rate Ratios (ITRRs) for the z114 and the previous-generation zSeries® processor families based upon measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user may experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated.

For more detailed performance information, consult the Large Systems Performance Reference (LSPR) available at

<https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lspindex>

Building a smarter business infrastructure

The IBM zEnterprise Unified Resource Manager

In today's IT infrastructure, silos of computing are not aligned with the businesses they support. Additionally, corporations struggle with regulatory compliance, volume of information, integrity, and security. These limitations, along with resource constraints and difficulties managing complexity and change, present increasing challenges in managing the IT infrastructure.

A smarter business infrastructure provides integrated visibility, control, and automation across all business and IT assets. The smarter business infrastructure is highly optimized to do more with less and facilitates the intelligent use of the information available. By managing and mitigating risks and utilizing flexible delivery choices, a smarter business infrastructure addresses the challenges facing business today.

Implementing a smarter business infrastructure is a strategy that can help to provide better visibility, control, and automation across the business, drive up levels of efficiency and optimization of both resources and people, bring order to the information explosion, and manage risk more productively. In addition, this type of strategy positions the infrastructure to leverage new types of delivery choices that can flex as needed, such as cloud computing, helping clients address today's most persistent business and IT constraints.

The IBM zEnterprise 114 provides a smarter business infrastructure through the introduction of the IBM zEnterprise Unified Resource Manager, which manages **System z ensembles** - collections of one or more zEnterprise System nodes in which each node is comprised of a z114 and its optionally attached IBM zEnterprise BladeCenter Extension (zBX) Model 002. An ensemble can consist of a single zEnterprise (z114 or a z196) with no zBX attached, using a fiber optic cable from an OSA-Express feature to an OSA-Express feature, or two-to-eight CPCs where at least one of the zEnterprises has a zBX attached. The resources of a zEnterprise System ensemble are managed and virtualized as a single pool of resources integrating system and workload management across the multisystem, multitier, multiarchitecture environment. Management actions for the ensemble are conducted from a designated primary Hardware Management Console. This HMC must be paired with an alternate HMC also designated to manage the ensemble, that can take over if the primary fails, in order to improve availability.

The functions of the Unified Resource Manager are grouped into two suites of tiered functionality that enable different levels of capability - **Manage suite** and **Automate suite** . Initial Manage suite functions are being delivered concurrently with the availability of the z114 with feature #0019. An enhanced set of Manage suite functions (#0019), as well as a set of Automate / Advanced Management firmware suite functions (#0020), are being delivered concurrently with the availability of the IBM zEnterprise BladeCenter Extension (zBX).

- The Manage suite is included in the base zEnterprise System at no charge for CPs, IFLs, zIIPs, and zAAPs.
- The Automate / Advanced Management firmware suite is included in the base zEnterprise System at no charge for CPs, zIIPs, and zAAPs. There is a charge for IFLs, blades, and IBM Smart Analytics Optimizer.

The Unified Resource Manager is designed to provide the following **initial Manage suite functions**:

- Monitoring and trend reporting of CPU energy efficiency, which can be helpful in managing the costs of deployed workloads
- Delivery of system activity using a new user interface monitors dashboard (augmenting the existing System Activity Display), enabling a broader and more granular view of system resource consumption

The Unified Resource Manager is designed to provide the following **enhanced Manage suite functions**:

- Integrated hardware management across all elements of the system, including operational controls, which is designed to deliver:
 - Licensed Internal Code (LIC) inventory, update, and service
 - Hardware and LIC problem detection, reporting, and call home
 - Field-guided repair and verification
 - Physical hardware configuration, backup, and restore
 - Primary/alternate replication and recovery for the HMC
 - Integrated hardware management, including operational controls, which provides a single point of control, a consistent interface, and comprehensive insight across the elements of the system to better manage operations, change control, and service
- Enhanced Monitoring and reporting of CPU energy efficiency with the ability to query maximum potential power.
- Fully automatic and coherent integrated resource discovery and inventory for all elements of the system without requiring user configuration, deployment of discovery libraries or sensors, or user scheduling of resource discovery. These integrated functions enable a simplified deployment methodology, decreasing complexity, and increasing accuracy and completeness for discovery and inventory.
- A private and physically isolated OSA-Express3 1000BASE-T Ethernet internal management network (the intranode management network - INMN), operating at 1 Gbps, that connects all zEnterprise System resources (CPCs, BladeCenters, blades, top-of-rack switches, power distribution units, and hypervisors) for management purposes. This INMN is pre-wired, internally switched, configured, and managed with full redundancy for high availability. These characteristics of the INMN enable simplified and secure management with no effect on data traffic. The INMN is exclusive to the IBM zEnterprise 114 and 196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.
- A private and secure OSA-Express3 or OSA-Express4S 10 Gigabit Ethernet intraensemble data network (IEDN) that connects all elements of a zEnterprise System ensemble. The IEDN is access-controlled using integrated virtual local area network (VLAN) provisioning. This requires no external user-managed switches or routers, which helps to reduce the need for firewalls and encryption, and simplifies network configuration and management, while providing full redundancy for high availability. IEDN management provides enforcement of strict access control across heterogeneous environments, further augmenting

security and simplicity. The intraensemble data network is exclusive to IBM zEnterprise 114 and 196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- Hypervisors, which, except for z/VM , are shipped, deployed, and serviced as zEnterprise System LIC and are booted automatically at power on reset. All hypervisors (PR/SM™ , z/VM , and PowerVM™) are isolated on the internal management network, providing intrusion prevention, integrity, secure virtual switches with integrated configuration, and monitoring, as well as problem management and reporting. Managed through the Unified Resource Manager, hypervisors become part of the trusted compute base and, except for z/VM , are deployed and maintained from a single point of control.
- Virtual server life-cycle management, which enables directed and dynamic virtual server provisioning across all hypervisors (PR/SM , z/VM , and PowerVM) from a single, uniform point of control. This enables integrated storage and network configuration and ensemble membership integrity. Virtual server life-cycle management provides an authoritative understanding of server, storage, and network elements within an ensemble. Augmenting the existing z/VM virtual server management functions available on System z10 , this support now also permits virtual servers to be created and deleted, and allows real and virtual networking resources to be managed by the Unified Resource Manager.

For more information on firmware managing the z/VM environment refer to Software Announcement [210-234](#), dated July 22, 2010 .

The Unified Resource Manager is designed to provide the following **Automate suite** functions:

- Representation of the physical and virtual resources that are used in the context of a deployed business function as a named workload. Workload representation enables focus on business goals rather than individual elements and better aligns resources with business needs.
- Ability to define a performance service level policy and to enable performance monitoring, reporting, and resource optimization that are aligned with user-defined workload service levels by allowing virtual CPU capacity to be adjusted automatically across virtual servers within a hypervisor. Automatic CPU adjustments are supported for z/VM and POWER7 blades.

POWER7 blades are supported in the zBX, providing logical device integration between System z and POWER7 for multitiered applications. POWER7 blades, along with PowerVM , are licensed separately and are enabled and managed, as part of the ensemble, by the Unified Resource Manager. For more information on POWER7 blades refer to " IBM BladeCenter P701 and P702 Express® blades deliver world's first IBM POWER7 processor-based blades to the marketplace," Hardware Announcement [110-061](#), dated April 13, 2010 .

IBM Smart Analytics Optimizer for DB2 for z/OS , V1.1 (5697-AQT) is being supported in the zBX, providing data-in-memory support to deliver significant performance and/or lower cost per transaction for queries redirected to the solution. It is enabled and managed, as part of the ensemble, by the Unified Resource Manager. For more information on IBM Smart Analytics Optimizer, refer to Software Announcement [210-266](#), dated July 22, 2010 , IBM Smart Analytics Optimizer for DB2 for z/OS , V1.1.

For more information on IBM zEnterprise BladeCenter Extension (zBX), refer to Hardware Announcement [110-177](#), dated July 22, 2010 , System z BladeCenter Extension.

Enhancements to IBM zEnterprise Unified Resource Manager

The IBM zEnterprise System continues to provide a smarter business infrastructure through augmentation of the IBM zEnterprise Unified Resource Manager. The Unified Resource Manager is designed to virtualize and manage all resources of a System z ensemble as a single pool of resources, integrating system and workload management across the multisystem, multitier, multiarchitecture environment.

Enhancements to the Manage suite (#0019) functions include :

- Hypervisor™ and virtual server life-cycle management to support the zBX integrated hypervisor for IBM System x blades (using Kernel-based Virtual Machine). Hypervisor and virtual server life-cycle management was first introduced with Unified Resource Manager for PR/SM , z/VM , and PowerVM .
- zEnterprise BladeCenter Extension Model 002 (zBX) integrated hypervisor for System x blades, which is shipped, deployed, and serviced as zEnterprise System Licensed Internal Code (LIC) and is booted automatically on each System x blade at power-on-reset. Like the existing supported hypervisors (PR/SM , z/VM , and PowerVM), it is isolated on the intranode management network (INMN) and managed by Unified Resource Manager, providing a single point of control and supporting directed and dynamic virtual server provisioning.

Enhancements to the Automate / Advanced Management firmware suite (#0020) functions include :

- Representation of the physical and virtual resources that are used in the context of a deployed business function as a named workload for IBM System x blades. Workload representation enables focus on business goals rather than individual elements and better aligns resources with business needs.
- Ability to define a performance service level policy and to enable performance monitoring and reporting for IBM System x blades.

System x blades running Linux on System x are supported in the zBX, utilizing the zBX integrated hypervisor for IBM System x blades (using Kernel-based Virtual Machine), providing logical device integration between System z and System x blades for multitiered applications. System x blades are licensed separately and are enabled and managed as part of the ensemble by the Unified Resource Manager.

For more information on IBM System x blades refer to Hardware Announcement [111-053](#), dated April 06, 2011 , IBM BladeCenter HX5 is a scalable blade server designed to provide new levels of utilization, performance, and reliability for compute- and memory-intensive workloads.

For more information regarding IBM zEnterprise BladeCenter Extension (zBX) Model 002, refer to Hardware Announcement [110-177](#), dated July 22, 2010 , IBM zEnterprise BladeCenter Extension (zBX).

Improved time coordination for zBX components: Network Time Protocol (NTP) clients, running on blades in the zBX, can synchronize their time to the NTP server provided by the Support Element (SE) every hour. Therefore, it is important for the SE's clock to maintain time accuracy.

An enhancement has been made to improve the time accuracy of the SE's Battery Operated Clock (BOC) by synchronizing the SE's BOC to the server's Time-of-Day (TOD) clock every hour, instead of the previous synchronization which took place every 24 hours. This enhancement allows the SE's clock to maintain a time accuracy of 100 milliseconds to an NTP server configured as the External Time Source in an STP-only CTN.

In addition, this enhancement provides the capability for the components in the zBX to maintain an approximate time accuracy of 100 milliseconds to an NTP server if they synchronize to the SE's NTP server at least once an hour. This enhancement is exclusive to the z196 and z114.

Ease of use enhancement

IPL from an alternate subchannel set

Two subchannel sets (MSS), were designed to provide greater I/O device configuration capabilities for large enterprises. Two subchannel sets could be defined to each channel subsystem - the base addresses and aliases in subchannel set 0 and the parallel access volume aliases in subchannel set 1; 63.75K subchannels in set-0 and the addition of 64K-1 subchannels in set-1. MSS was delivered on the System

z9 in September of 2005 and supported by z/OS V1.7 and Linux on System z . This applied to the ESCON and FICON protocols.

Now, we are introducing IPL from an alternate subchannel set to allow enterprises to IPL from subchannel set 1 (SS1) in addition to subchannel set 0. Devices used early during IPL processing can now be accessed using subchannel set 1 This is intended to allow the users of Metro Mirror (PPRC) secondary devices defined using the same device number and a new device type in an alternate subchannel set to be used for IPL, IODF, and stand-alone dump volumes when needed.

IPL from an alternate subchannel set is exclusive to the z196 and z114, and applies to the FICON and zHPF protocols (CHPID type FC), and is supported by z/OS . Refer to the [Software requirements](#) section. It is applicable to all of the supported FICON features.

Reliability and availability

With the introduction of the IBM zEnterprise 114, IBM continues to deliver enterprise reliability and availability with enhancements to the memory subsystem.

Major redesign of memory subsystem for improved availability

IBM's most robust error correction to date can be found in the memory subsystem. A new redundant array of independent memory (RAIM) technology is being introduced to provide protection at the dynamic random access memory (DRAM), dual inline memory module (DIMM), and memory channel level. Three full DRAM failures per rank can be corrected. DIMM-level failures, including components such as the controller application specific integrated circuit (ASIC), the power regulators, the clocks, and the board, can be corrected. Memory channel failures such as signal lines and control lines can be corrected. Upstream and downstream data signals can be spared using two spare wires on both the upstream and downstream paths. One of these signals can be used to spare a clock signal line (one up stream and one down stream). Together these improvements are designed to deliver System z's most resilient memory subsystem to date.

Cryptography enhancements

CP Assist for Cryptographic Function (CPACF) enhancements

The following exploit Message-Security-Assist Extension 4:

New instructions:

- Cipher Message with CFB (KMF)
- Cipher Message with Counter (KMCTR)
- Cipher Message with OFB (KMO)

New function code for existing instruction:

- Compute Intermediate Message Digest (KIMD), and extension for GHASH

More information on CPACF can be found in " IBM System z10 - Delivering security-rich offerings to protect your data," Hardware Announcement [109-678](#), dated October 20, 2009 .

This Crypto function is exclusive to IBM zEnterprise 114 and 196 servers.

Common Cryptographic Architecture and Crypto Express3 enhancements

To continue to satisfy user requirements, new functions are now available for the Common Cryptographic Architecture (CCA) and the Crypto Express3 feature, when defined as a coprocessor, complementing the Central Processor Assist for

Cryptographic Function (CPACF). The CPACF provides high-performance hardware encryption and decryption on every z196 and z114 processor unit defined as an IFL.

Crypto Express3 remains a tamper-sensing and tamper-responding programmable cryptographic feature. Each Crypto Express3 PCIe adapter contains dual processors that operate in parallel to support the Common Cryptographic Architecture with high reliability. The Crypto Express3 feature continues to reside in the I/O cage or I/O drawer of the zEnterprise System.

Common Cryptographic Architecture enhancements

The following enhancements have been added to the Common Cryptographic Architecture support which is used in the Crypto Express3 feature when it is configured as a coprocessor.

Crypto Express3 enhancements

Crypto Express3 represents the newest-generation cryptographic feature designed to complement the cryptographic functions of CP Assist for Cryptographic Function (CPACF). The Crypto Express3 feature, residing in the I/O cage or I/O drawer of the z114, continues to support all of the cryptographic functions available on Crypto Express3 on z10 .

For more information on Crypto Express3 refer to " IBM System z10 - Delivering security-rich offerings to protect your data," Hardware Announcement [109-678](#), dated October 20, 2009 .

Enhanced Driver Maintenance (EDM) and Concurrent Machine Change Level (MCL) apply: a process to reduce outages for new releases. With Enhanced Driver Maintenance and Concurrent MCL apply, most new cryptographic functions can be applied without cryptographic coprocessor card off / on. It is now possible to upgrade Common Cryptographic Architecture (CCA), segment 3, Licensed Internal Code without any performance impact during the upgrade. However, some levels of CCA or hardware changes still require cryptographic coprocessor card vary off / on.

This Crypto function is exclusive to IBM zEnterprise z114 and z196 servers.

When one or both of the two PCIe adapters of a Crypto Express3 feature are configured as a coprocessor, the following cryptographic enhancements are supported:

- **ANSI X9.8 PIN security** , which facilitates compliance with the processing requirements defined in the new version of the ANSI X9.8 and ISO 9564 PIN Security Standards and provides added security for transactions that require Personal Identification Numbers (PINs).

This Crypto function is exclusive to IBM zEnterprise z114 and z196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **Enhanced Common Cryptographic Architecture (CCA)** , a key wrapping to comply with ANSI X9.24-1 key bundling requirements, is a Common Cryptographic Architecture (CCA) key token wrapping method using Cipher Block Chaining (CBC) mode in combination with other techniques to satisfy the key bundle compliance requirements in standards including ANSI X9.24-1 and the recently published Payment Card Industry Hardware Security Module (PCI HSM) standard.

This Crypto function is exclusive to IBM zEnterprise 114 and 196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **Secure Keyed-Hash Message Authentication Code (HMAC)**, a method for computing a message authentication code using a secret key and a secure hash function. HMAC is defined in the standard FIPS (Federal Information Processing Standard) 198, "The Keyed-Hash Message Authentication Code." The new CCA functions support HMAC using SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 hash algorithms. The HMAC keys are variable-length and are securely encrypted so that their values are protected.

This Crypto function is exclusive to IBM zEnterprise 114 and 196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **Elliptical Curve Cryptography Digital Signature Algorithm**, an emerging public-key algorithm to eventually replace RSA cryptography in many applications. ECC is capable of providing digital signature functions and key agreement functions. The new CCA functions provide ECC key generation and key management and provide digital signature generation and verification functions compliance with the ECDSA method described in ANSI X9.62 "Public Key Cryptography for the Financial Services Industry: The Elliptical Curve Digital Signature Algorithm (ECDSA)." ECC uses keys that are shorter than RSA keys for equivalent strength-per-key-bit. With RSA impractical at key lengths with strength-per-key-bit equivalent to AES-192 and AES-256, the strength-per-key-bit is substantially greater in an algorithm that uses elliptic curves.

This Crypto function is exclusive to IBM zEnterprise z114 and z196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

When one or both of the two PCIe adapters of a Crypto Express3 feature are configured as an accelerator, the following cryptographic enhancement is supported:

- **Modulus Exponent (ME) and Chinese Remainder Theorem (CRT)**, RSA encryption and decryption with key lengths greater than 2048 bits and up to 4096 bits.

This Crypto function is exclusive to IBM zEnterprise z114 and z196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **Expanded key support for AES algorithm** : CCA currently supports the Advanced Encryption Standard (AES) algorithm to allow the use of AES keys to encrypt data. Expanded key support for AES adds a framework to support a much broader range of application areas, and lays the groundwork for future use of AES in areas where standards and user applications are expected to evolve.

As stronger algorithms and longer keys become increasingly common, security requirements dictate that these keys must be wrapped using key encrypting keys (KEKs) of sufficient strength. This feature adds support for AES key encrypting keys. These AES wrapping keys have adequate strength to protect other AES keys for transport or storage. The new AES key types use the variable-length key token. The supported key types are EXPORTER, IMPORTER, and for use in the encryption and decryption services, CIPHER. New APIs have been added or modified to manage and use these new keys.

The following new or modified CCA API functions are also supported:

- Key Token Build2: Builds skeleton variable-length key tokens
- Key Generate2: Generates keys using random key data
- Key Part Import2: Creates keys from key part information
- Key Test2: Verifies the value of a key or key part
- Key Translate2
 - Translates a key: Changes the key encrypting key (KEK) used to wrap a key
 - Reformats a key: Converts keys between the legacy token format and the newer variable-length token format
- Symmetric Key Export: Modified to also export AES keys
- Symmetric Key Import2: Imports a key that has been wrapped in the new token format
- Secure Key Import2: Wraps key material under the master key or an AES KEK
- Restrict Key Attribute: Changes the attributes of a key token
- Key Token Parse2: Parses key attributes in the new key token
- Symmetric Algorithm Encipher and Symmetric Algorithm Decipher: Enhanced to encipher and decipher data using AES keys wrapped in the new variable-length token format

Expanded key support for AES algorithm and CCA API functions supported by the Crypto Express3 feature when defined as a coprocessor is exclusive to the z196 and z114, and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **Enhanced ANSI TR-31 interoperable secure key exchange** : ANSI TR-31 defines a method of cryptographically protecting Triple Data Encryption Standard (TDES) cryptographic keys and their associated usage attributes. The TR-31 method complies with the security requirements of the ANSI X9.24 Part 1 standard, although use of TR-31 is not required in order to comply with that standard. CCA has added functions that can be used to import and export CCA TDES keys in TR-31 formats. These functions are designed primarily as a secure method of wrapping TDES keys for improved and more secure key interchange between CCA and non-CCA devices and systems.

Enhanced ANSI TR-31 interoperable secure key exchange supported by the Crypto Express3 feature when defined as a coprocessor is exclusive to the z196 and z114, and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **PIN block decimalization table protection** : To help avoid a decimalization table attack to learn a Personal Identification Number (PIN), a solution is now available in the CCA API to thwart this attack by protecting the decimalization table from manipulation. PINs are most often used for automated teller machines (ATMs) but are increasingly used at point-of sale, for debit and credit cards.

PIN block decimalization table protection supported by the Crypto Express3 feature when defined as a coprocessor is exclusive to the z196 and z114, and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **PKA RSA OAEP with SHA-256 algorithm** : RSA Encryption Scheme - Optimal Asymmetric Encryption Padding (RSA OAEP) is a public-key encryption scheme or method of encoding messages and data in combination with the RSA algorithm and a hash algorithm.

Currently, the Common Cryptographic Architecture and z/OS Integrated Cryptographic Service Facility (ICSF) provide key management services supporting the RSA OAEP method using the SHA-1 hash algorithm, as defined by the public key cryptographic standards (PKCS) #1 V2.0 standard. These services can be used to exchange AES or DES/TDES key values securely between financial institutions and systems. However, PKCS#1 V2.1 extends the OAEP method to include the use of the SHA-256 hashing algorithm to increase the strength of the key wrapping and unwrapping mechanism. The CCA key management services have been enhanced so that they can use RSA OAEP with SHA-256 in addition to RSA OAEP with SHA-1.

This provides support for PKCS that is mandated by some countries for interbank transactions and communication systems.

PKA RSA OAEP with SHA-256 algorithm is supported by the Crypto Express3 feature when defined as a coprocessor, is exclusive to z196 and z114, and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

- **Additional Elliptic Curve Cryptography (ECC) functions** : The Common Cryptographic Architecture has been extended to include the Elliptic Curve Diffie-Hellman (ECDH) algorithm.

ECDH is a key agreement protocol that allows two parties, each having an elliptic curve public-private key pair, to establish a shared secret over an insecure channel. This shared secret may be used directly as a key, or to derive another key which can then be used to encrypt subsequent communications using a symmetric key cipher such as AES KEK.

Enhancements include:

- Updated key management function to support AES KEKs
 - Generating an ECC private key wrapped with an AES KEK

- Importing and exporting an ECC private key wrapped with an AES KEK
- Support for ECDH with a new service

These additional Elliptic Curve Cryptographic functions are supported by the Crypto Express3 feature when defined as a coprocessor, are exclusive to z196 and z114, and are supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

When one or both of the two PCIe adapters of a Crypto Express3 feature are configured as an accelerator, the following cryptographic enhancement is supported:

- **Modulus Exponent (ME) and Chinese Remainder Theorem (CRT)**, RSA encryption and decryption with key lengths greater than 2048-bits and up to 4096-bits.

This Crypto function is exclusive to IBM zEnterprise z114 and z196 servers and is supported by z/OS and z/VM . Refer to the [Software requirements](#) section.

Trusted Key Entry 7.1 Licensed Internal Code (LIC)

The following functions are supported in the TKE 7.1 LIC:

- **New access control support for all TKE applications** : Every TKE application and the ability to create and manage crypto module and domain groups now require the TKE local cryptographic adapter profile to have explicit access to the TKE application or function the user wants to run. This was done to provide more control of what functions TKE users are allowed to perform.
- **New migration utility** : During a migration from a lower release of TKE to TKE 7.1 LIC, it will be necessary to add access control points to the existing roles. The new access control points can be added through the new Migrate Roles Utility or by manually updating each role through the Cryptographic Node Management Utility. The IBM-supplied roles created for TKE 7.1 LIC have all of the access control points needed to perform the functions they were permitted to use in TKE releases prior to TKE 7.1 LIC.
- **Single process for loading an entire key** : The TKE now has a wizard-like feature that takes users through the entire key loading procedure for a master or operational key. The feature preserves all of the existing separation of duties and authority requirements for clearing, loading key parts, and completing a key. The procedure saves time, by walking users through the key loading procedure. However, this feature does not reduce the number of people it takes to perform the key load procedure.
- **Single process for generating multiple key parts of the same type** : The TKE now has a wizard-like feature that allows a user to generate more than one key part at a time. The procedure saves time because the user has to start the process only one time, and the TKE efficiently generates the desired number of key parts.
- **AES operational key support** : CCA V4.2 for the Crypto Express3 feature includes three new AES operational key types. From the TKE, users can load and manage the new AES EXPORTER, IMPORTER, and CIPHER operational keys from the TKE workstation crypto module notebook.
- **Decimalization table support** : CCA V4.2 for the Crypto Express3 feature includes support for 100 decimalization tables for each domain on a Crypto Express3 feature. From the TKE, users can manage the decimalization tables on the Crypto Express3 feature from the TKE workstation crypto module notebook. Users can manage the tables for a specific domain or manage the tables of a set of domains if they are using the TKE workstation Domain Grouping function.
- **Host cryptographic module status support** : From the TKE workstation crypto module notebook, users will be able to display the current status of the host cryptographic module that is being managed. If they view the Crypto Express3 feature module information from a crypto module group or a domain group, they will see only the status of the group's master module.
- **Display of active IDs on the TKE console** : A user can be logged onto the TKE workstation in privileged access mode. In addition, the user can be signed

onto the TKE workstation's local cryptographic adapter. If a user is signed on in privileged access mode, that ID is shown on the TKE console. With this new support, both the privileged access mode ID and the TKE local cryptographic adapter ID will be displayed on the TKE console.

- **Increased number of key parts on a smart card** : If a TKE smart card is initialized on a TKE workstation with a 7.1 level of LIC, it will be able to hold up to 50 key parts. Previously, TKE smart cards could hold only 10 key parts.
- **Use of ECDH to derive the shared secret** : When the TKE workstation with a 7.1 level of LIC exchanges encrypted material with a Crypto Express3 at CCA level V4.2, Elliptic Curve Diffie-Hellman (ECDH) is used to derive the shared secret. This increases the strength of the transport key used to encrypt the material.

Trusted Key Entry 7.1 Licensed Internal Code is supported by the Crypto Express3 feature when defined as a coprocessor on z196, z114, z10 EC, and z10 BC servers. TKE workstation (#0841) is required.

I/O improvements

Assigning WWPNS to physical Fibre Channel Protocol (FCP) ports

This support extends the capabilities of the worldwide port name (WWPN) prediction tool to **physical** Fibre Channel Protocol (FCP) channel/ports, allowing the tool to now show WWPNS for both virtual and physical ports ahead of system installation time. In addition, you can retain your storage area network (SAN) configuration if a FICON feature is replaced instead of altering your SAN configuration based upon the FICON feature's "burned in" WWPN.

This enhancement applies to all of the FICON features supported on z114 when configured as CHPID type FCP, supporting attachment to SCSI devices. It is transparent to operating systems. For an update to the WWPN prediction tool, refer to the *Tools* section of Resource Link® at planned availability of the z114.

<http://www.ibm.com/servers/resourcelink>

z/OS discovery and autoconfiguration for FICON channels : With z114 and z/OS , a new function, z/OS discovery and autoconfiguration (zDAC), is designed to automatically perform a number of I/O configuration definition tasks for new and changed disk and tape controllers connected to a switch or director when attached to a FICON channel. When new controllers are added to an I/O configuration or changes are made to existing controllers, the system is designed to discover them and propose configuration changes based on a policy you define in the hardware configuration dialog (HCD). Your policy can include preferences for availability and bandwidth including parallel access volume (PAV) definitions, control unit numbers, and device number ranges.

zDAC is designed to perform discovery for all systems in a sysplex that support the function. The proposed configuration will incorporate the current contents of the I/O definition file (IODF) with additions for newly installed and changed control units and devices. zDAC is designed to help simplify I/O configuration on the z114 running z/OS and reduce complexity and setup time.

zDAC applies to all FICON features supported on z114 when configured as CHPID type FC and is supported by z/OS . zDAC is exploited by the IBM System Storage® DS8700. For more information refer to Hardware Announcement [110-168](#), dated July 20, 2010 , IBM System Storage DS8700 (M/T 242x) delivers z/OS Distributed Data Backup and new functional capabilities. Refer to the [Software requirements](#) section.

PCIe I/O drawer - increased capacity, granularity, bandwidth, and RAS

A new I/O drawer and new form factor I/O cards are being introduced to support a direct Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) infrastructure with increased capacity, granularity, and infrastructure bandwidth, as well as increased reliability, availability, and serviceability. New hardware features

are being introduced for the SAN and the LAN to support PCIe Gen2 - FICON Express8S and OSA-Express4S.

Capacity: With its compact design, the PCIe I/O drawer occupies the same space as one I/O drawer while delivering 100% more capacity - up to 64 channels (32 I/O cards) are supported versus the 32 channels (eight I/O cards) offered with the I/O drawer.

Granularity : With the PCIe I/O drawer, IBM is introducing enhanced granularity features for the storage area network (SAN) and the local area network (LAN).

- For the FICON/zHPF/FCP storage area network: FICON Express8S for single mode fiber (LX) and multimode fiber (SX) environments with two channels per feature, and two channel path identifiers (CHPIDs).
- For the 10 Gigabit Ethernet and Gigabit Ethernet local area network: OSA-Express4S for single mode fiber (LX, LR) and multimode fiber (SX, LR) environments. The 10 GbE features have one port per feature and one CHPID. The GbE features have two ports per feature and one CHPID shared by the two ports.

Increased infrastructure bandwidth : The new I/O infrastructure introduces the industry-standard Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) infrastructure from the processor book to the I/O cards, which now supports an 8 gigabytes per second (8 GBps) bus. This new infrastructure offers enhanced bandwidth in the connection to the processor book as well as the distribution buses contained in the drawer.

Reliability, Availability, Serviceability : The PCIe I/O drawer is designed with symmetrical, redundant cooling across all I/O cards and power supplies for improved RAS. Dual PCIe paths to the system processors from an I/O card provide enhanced resiliency. Also included are temperature monitoring of critical ASICs to optimize cooling and reliability.

FICON Express8S - a new generation for FICON , zHPF, and FCP

A new generation of features for the storage area network (SAN) is being introduced in support of the PCIe 8 GBps host bus and the PCIe I/O drawer. The new features for the multimode and single mode fiber optic cabling environments have path length reductions for High Performance FICON for System z (zHPF) and Fibre Channel Protocol (FCP), increased start I/Os, improved throughput for zHPF and FCP with the introduction of a hardware data router, and increased port granularity - two channels/ports per feature.

New design for increased performance for zHPF and FCP : FICON Express8S contains a new IBM ASIC which is designed to support the 8 GBps PCIe interface to the PCIe I/O drawer and increased start I/Os. In addition, a hardware data router has been added in support of the zHPF and FCP protocols for path length reduction and increased throughput. FICON Express8S supports a link data rate of 2, 4, or 8 Gbps autonegotiated.

With these changes FICON Express8S, when supporting the zHPF or FCP protocols, has been designed to achieve full duplex line speed - 8 Gbps - in each direction.

Increased port granularity : The FICON Express8S 10KM LX and SX features for single mode and multimode fiber optic cabling environments each now have two channels/ports per feature versus the four channels per feature for the FICON Express8 features. This design helps facilitate purchasing the right number of ports to help satisfy your application requirements and to better optimize for redundancy.

The FICON Express8S features, supporting CHPID types FC (zHPF, FICON , channel-to-channel) and FCP, are exclusive to the z196 and z114. They are for use exclusively in the PCIe I/O drawer and are supported by z/OS , z/VM , z/VSE , z/TPF, and Linux on System z . Refer to the [Software requirements](#) section.

The High Performance FICON for System z journey:

- In October of 2008 IBM made the initial zHPF announcement on System z10 with FICON Express4 and FICON Express2 channels.
 - Maximum of 31k zHPF IOs/sec, 2.2 times the FICON protocol
 - Single track limit for zHPF data transfers
- In July 2009 IBM introduced FICON Express8 channels on System z10 .
 - Maximum of 52k zHPF IOs/sec
 - 64k byte limit for zHPF data transfers
- In July 2010 IBM introduced additional support on zEnterprise 196. This support is being carried forward to the z114.
 - Extension to multitracks of zHPF data transfers
 - Support for multiple tracks of data
 - Removal of 64 k byte data transfer limit
 - The ability of channels to operate at rates designed to fully exploit the bandwidth of a FICON Express8 channel
- Today IBM introduces the new FICON Express8S channel.
 - A hardware data router is introduced for more efficient zHPF data transfers.
 - The FICON Express8S channel is the first channel with hardware specifically designed to support zHPF.
 - FICON Express8, FICON Express4, and FICON Express2 have a firmware-only zHPF implementation.

In laboratory measurements, using FICON Express8S in a zEnterprise with the zHPF protocol and small data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum of 92,000 IOs/sec, compared to the maximum of 52,000 IOs/sec achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 77% increase and applies to reads, writes, and a read/write mix.

In laboratory measurements, using FICON Express8S in a zEnterprise with the zHPF protocol and a mix of large sequential read and write data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum throughput of 1600 MB/sec (reads + writes) compared to a maximum of 770 MB/sec (reads + writes) achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 108% increase.

This performance data was measured in a controlled environment running an I/O driver program under z/OS . The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

Increased performance for the FCP protocol : A FICON Express8S feature, when defined as CHPID type FCP, conforms to the Fibre Channel Protocol (FCP) standard to support attachment of SCSI devices, to complement the classical storage attachment supported by FICON and zHPF channels.

In laboratory measurements, using FICON Express8S in a zEnterprise with the FCP protocol for small data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum of 92,000 IOs/sec, compared to the maximum of 84,000 IOs/sec achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 10% increase and applies to reads, writes, and a read/write mix.

In laboratory measurements, using FICON Express8S in a zEnterprise with the FCP protocol and an internal driver supporting the hardware data router, executing a mix of large sequential read and write data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum throughput of 1600 MB/sec (reads + writes) compared to the maximum of 770 MB/sec (reads + writes) achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 108% increase.

The FCP protocol is supported by z/VM , z/VSE , and Linux on System z . Refer to the [Software requirements](#) section.

FCP channels to support T10-DIF for enhanced reliability

Recognizing that high reliability is important to maintaining the availability of business-critical applications, the System z Fibre Channel Protocol (FCP) has implemented support of the American National Standards Institute's (ANSI) T10 Data Integrity Field (DIF) standard. Data integrity protection fields are generated by the operating system and propagated through the storage area network (SAN). System z helps to provide added end-to-end data protection between the operating system and the storage device.

An extension to the standard, Data Integrity Extensions (DIX), provides checksum protection from the application layer through the host bus adapter (HBA), where cyclical redundancy checking (CRC) protection is implemented.

T10-DIF support by the FICON Express8S and FICON Express8 features, when defined as CHPID type FCP, is exclusive to z196 and z114. Exploitation of the T10-DIF standard requires support by the operating system and the storage device. Refer to the [Software requirements](#) section.

Increased port granularity: The FICON Express8S 10KM LX and SX features for single mode and multimode fiber optic cabling environments each now have two channels/ports per feature versus the four channels per feature for the FICON Express8 features. This design helps facilitate purchasing the right number of ports to help satisfy your application requirements and to better optimize for redundancy.

The FICON Express8S features, supporting CHPID types FC (zHPF, FICON , channel-to-channel) and FCP, are exclusive to z196 and z114. They are for use exclusively in the PCIe I/O drawer and are supported by z/OS , z/VM , z/VSE , z/TPF, and Linux on System z . Refer to the [Software requirements](#) section.

Access to high-speed intraserver network - HiperSockets™

A single logical partition can connect up to 32 HiperSockets , double the number of HiperSockets previously supported. With HiperSockets , you have independent, integrated, virtual local area networks (LANs). No physical cabling or external connections are required. HiperSockets supports Internet Protocol version 4 (IPv4) and IPv6.

Up to 32 HiperSockets are supported by z/OS , z/VM , z/VSE, and Linux on System z . Refer to the [Software requirements](#) section.

HiperSockets network traffic analyzer (HS NTA) for Linux on System z environments

An enhancement to the HiperSockets architecture is designed to help simplify problem isolation and resolution. You now have the capability to trace Layer 2 and Layer 3 HiperSockets network traffic.

HS NTA allows Linux on System z to control the trace for the internal virtual LAN, to capture the records into host memory and storage (file systems). Linux on System z tools can be used to format, edit, and process the trace records for analysis by system programmers and network administrators. Refer to the [Software requirements](#) section.

HiperSockets network traffic analyzer is supported on IBM zEnterprise and IBM System z10 servers.

Networking enhancements

Access to the LAN with OSA-Express3

The OSA-Express3 family of local area network (LAN) adapters is designed for use in high-speed enterprise backbones, for local area network connectivity between campuses, to connect server farms to the mainframe, and to consolidate file servers onto the mainframe.

The OSA-Express3 features incorporated a hardware data router in the design, as well as a faster microprocessor and PCIe bus; functions and components that help to reduce latency and improve throughput. Choose the features that satisfy your infrastructure requirements, now and in the future.

OSA-Express3 Ethernet features:

- 10 Gigabit Ethernet long reach (10 GbE LR) for single mode fiber infrastructures; supports 10 Gbps link data rate
- 10 Gigabit Ethernet short reach (10 GbE SR) for multimode fiber infrastructures; supports 10 Gbps link data rate
- Gigabit Ethernet long wavelength (GbE LX) for single mode fiber infrastructures; supports 1 Gbps link data rate
- Gigabit Ethernet short wavelength (GbE SX) for multimode fiber infrastructures; supports 1 Gbps link data rate
- 1000BASE-T Ethernet for Category 5 or Category 6 copper infrastructures; autonegotiates to 10, 100, or 1000 Mbps

Note: The four OSA-Express3 features (excluding the 1000BASE-T feature) listed above are not orderable on the IBM zEnterprise 114 models. If they are installed at the time of an upgrade to the IBM zEnterprise 114 they may be retained. OSA-Express4S is the primary OSA offering. The OSA-Express3 Ethernet features support:

- Queued Direct Input/Output (QDIO) - uses memory queues and a signaling protocol to directly exchange data between the OSA microprocessor and the network software for high-speed communication.
 - QDIO Layer 2 (Link layer) - for IP (IPv4, IPv6) or non-IP (AppleTalk DECnet, IPX, NetBIOS, or SNA) workloads. Using this mode the Open Systems Adapter (OSA) is protocol-independent and Layer-3 independent. Packet forwarding decisions are based upon the Medium Access Control (MAC) address.
 - QDIO Layer 3 (Network or IP layer) - for IP workloads. Support for packet forwarding is based upon IP address or virtual MAC address. Guests may share OSA's MAC address or guests may have their own unique virtual MAC address.

| CHPID type | Applicable features | Purpose/Traffic |
|------------|---------------------|--|
| OSC | 1000BASE-T | OSA-Integrated Console Controller (OSA-ICC). Supports TN3270E, non-SNA DFT to IPL CPCs & LPs. |
| OSD | All | Queue Direct Input/Output (QDIO) architecture; TCP/IP traffic when Layer 3 (uses IP address) and Protocol-independent when Layer 2 (uses MAC address). |
| OSE | 1000BASE-T | Non-QDIO; For SNA/APPN/HPR traffic and TCP/IP passthru traffic. |
| OSM | 1000BASE-T | OSA-Express for Unified Resource Manager. Connectivity to intranode management network (INMN) from z114 to Unified Resource Manager functions. |
| OSN | GbE 1000BASE-T | OSA-Express for NCP; Appears to OS as a device supporting channel data link control (CDLC) protocol. Enables Network Control Program (NCP) channel-related functions such as loading and dumping to NCP. Provides LP-to-LP connectivity OS to IBM Communication Controller for Linux (CCL). OSN CHPIDs are not available on the OSA-Express4S. |
| OSX | 10 GbE | OSA-Express for ZBX. Provides connectivity |

and access control to the intraensemble data network (IEDN) from z114 to Unified Resource Manager functions.

For network performance - inbound workload queuing :

z/OS workloads are becoming more diverse, and each type of workload may have unique service level requirements. Starting with OSA-Express3, inbound workload queuing (IWQ) creates multiple input queues and allows OSA to differentiate workloads "off the wire" and then assign work to a specific input queue (per device) to z/OS . With each input queue representing a unique type of workload, each having unique service and processing requirements, the IWQ function allows z/OS to preassign the appropriate processing resources for each input queue. This approach allows multiple concurrent z/OS processing threads to then process each unique input queue (workload), avoiding traditional resource contention. In a heavily mixed workload environment, this "off the wire" network traffic separation provided by OSA-Express3 IWQ reduces the conventional z/OS processing required to identify and separate unique workloads, which results in improved overall system performance and scalability.

A primary objective of IWQ is to provide improved performance for business-critical interactive workloads by reducing contention created by other types of workloads. The types of z/OS workloads that are identified and assigned to unique input queues are:

1. z/OS Sysplex Distributor traffic - Network traffic that is associated with a distributed dynamic virtual IP address (DVIPA) is assigned a unique input queue, allowing the Sysplex Distributor traffic to be immediately distributed to the target host.
2. z/OS bulk data traffic - Network traffic that is dynamically associated with a streaming (bulk data) TCP connection is assigned to a unique input queue, allowing the bulk data processing to be assigned the appropriate resources and isolated from critical interactive workloads.
3. Enterprise Extender traffic - Inbound workload queuing (IWQ) for the OSA-Express features has been enhanced to differentiate and separate inbound Enterprise Extender traffic to a new input queue. The Enterprise Extender separation and processing associated with the Enterprise Extender input queue provides improved scalability and performance for Enterprise Extender.

IWQ is supported starting on the z10 , and is available starting with OSA-Express3 CHPID types OSD and OSX (exclusive to the z114 and z196). IWQ is also supported by the z/OS operating system and by z/VM for guests. Refer to the [Software requirements](#) section.

For network management - query and display OSA configuration

Previously, OSA-Express system architecture introduced the capability for operating systems to dynamically register the OSA configuration. This approach significantly improved the OSA-Express usability by reducing the burden placed on the system administrator to manually configure OSA-Express for each unique operating system configuration. Traditionally, the Open Systems Adapter Support Facility (OSA/SF) has provided the administrator with the ability to monitor the OSA configuration.

As additional complex functions have been added to OSA, the ability for the system administrator to display, monitor, and verify the specific current OSA configuration unique to each operating system has become more complex. OSA-Express3 introduces the capability for the operating system to directly query and display the current OSA configuration information (similar to OSA/SF). z/OS exploits this new OSA capability by introducing a new TCP/IP operator command called **Display OSAINFO** . Display OSAINFO allows the operator to monitor and verify the current OSA configuration, which helps to improve the overall management, serviceability, and usability of OSA-Express3.

Display OSAINFO is available starting with OSA-Express3 supporting CHPID types OSD, OSM, and OSX, the z/OS operating system, and z/VM for guest exploitation. Refer to the [Software requirements](#) section.

OSA-Express4S - a new form factor of Ethernet for the LAN

A new generation of 10 Gigabit Ethernet (10 GbE) and Gigabit Ethernet (GbE) OSA-Express features is being introduced in support of the PCIe 8 GBps host bus and the PCIe I/O drawer. These new features retain the robust performance characteristics of the OSA-Express3 GbE and 10 GbE features while providing increased port granularity.

Port granularity for increased flexibility : The OSA-Express4S features have half the number of ports per feature compared to OSA-Express3. This design facilitates purchasing the right number of ports to help satisfy your application requirements and to better optimize for redundancy.

OSA-Express4S 10 GbE with one port per feature : The OSA-Express4S 10 GbE features have one CHPID per feature and one port associated with the CHPID. A feature can be configured to support CHPID type OSD or OSX.

- CHPID type OSD continues to support the Queued Direct Input/Output (QDIO) architecture, TCP/IP traffic when Layer 3 (uses IP address), protocol-independent when Layer 2 (uses MAC address).
- CHPID type OSX continues to support connectivity and access control to the intraensemble data network (IEDN) from the z196 or z114 to zBX (OSA-Express 10 GbE features only)

OSA-Express4S 10 Gigabit Ethernet Long Reach (LR) is available for single mode fiber environments. OSA-Express4S 10 Gigabit Ethernet Short Reach (SR) is available for multimode fiber environments.

OSA-Express4S GbE with two ports per feature : The OSA-Express4S GbE features have one CHPID per feature (CHPID type OSD) and two ports associated with the CHPID.

Note: CHPID type OSN (OSA-Express for NCP) is not being offered for the OSA-Express4S GbE features.

OSA-Express4S GbE long wavelength (LX) is available for single mode fiber environments. OSA-Express4S GbE short wavelength (SX) is available for multimode fiber environments.

The OSA-Express4S 10 Gigabit Ethernet and Gigabit Ethernet features are exclusive to the z196 and z114. They are for use exclusively in the PCIe I/O drawer and are supported by z/OS , z/VM , z/VSE , z/TPF, and Linux on System z . Refer to the [Software requirements](#) section.

Checksum offload for IPv6 packets is now available for z/OS environments : When the checksum function is offloaded from the host, CPU cycles are reduced, improving performance. With the introduction of OSA-Express4S, the checksum offload function is now performed for IPv6 packets as well as IPv4 packets, whether the traffic goes out to the local area network (LAN), comes in from the LAN, or flows logical partition-to-logical partition through OSA-Express4S.

Checksum offload provides the capability of calculating the Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Protocol (IP) header checksums for Internet Protocol Version 4 (IPv4) packets and now, IPv6 packets.

When checksum offload was introduced in May of 2003, it was limited to IPv4 packets. Checksum offload for IPv6 packets is exclusive to OSA-Express4S features (CHPID types OSD and OSX) on the z196 and z114. It is supported by z/OS . Refer to the [Software requirements](#) section. Checksum offload for IPv4 packets is currently available for all in-service releases of z/OS and Linux on System z .

Checksum offload for LPAR-to-LPAR traffic in the z/OS environment is included in the OSA-Express4S design for both IPv4 and IPv6 packets. Refer to the [Software requirements](#) section.

Large send for IPv6 packets: Large send (also referred to as TCP segmentation offload) is designed to improve performance by offloading outbound TCP segmentation processing from the host to an OSA-Express4S feature by employing a more efficient memory transfer into OSA-Express4S.

Large send support for IPv6 packets applies to the OSA-Express4S features (CHPID type OSD and OSX), and is exclusive to the z196 and z114. Refer to the [Software requirements](#) section. Large send for IPv4 packets is currently available on OSA-Express2, OSA-Express3, and OSA-Express4S for all in-service releases of z/OS , Linux on System z , and z/VM for guest exploitation.

Note: Large Send is not supported for LPAR-to-LPAR packets.

Inbound Workload Queuing for Enterprise Extender - for improved scalability and performance : Inbound workload queuing (IWQ) for the OSA-Express features has been enhanced to differentiate and separate inbound Enterprise Extender traffic to a new input queue. The Enterprise Extender separation and processing associated with the Enterprise Extender input queue provides improved scalability and performance for Enterprise Extender.

With each input queue representing a unique type of workload, each having unique service and processing requirements, the IWQ function allows z/OS to use appropriate processing resources for each input queue. This approach allows multiple concurrent z/OS processing threads to process each unique input queue to avoid traditional resource contention. In a heavily mixed workload environment, this "off the wire" network traffic separation provided by OSA-Express IWQ reduces the conventional z/OS processing required to identify and separate unique workloads.

Inbound workload queuing for Enterprise Extender is supported by the OSA-Express4S and OSA-Express3 features when defined as CHPID type OSD or OSX. It is supported by z/OS and by z/VM for guest exploitation. Refer to the [Software requirements](#) section.

For more information on inbound workload queuing, refer to Hardware Announcement [110-170](#), dated July 22, 2010 , The IBM zEnterprise System -- A new dimension in computing.

Design change yields performance benefit for OSA-Express4S 10 Gigabit Ethernet inbound traffic

Performance using jumbo frames: In laboratory measurements, using an OSA-Express4S 10 Gigabit Ethernet (10 GbE) feature in a z196 defined as CHPID type OSD with an inbound-to-the-host streams workload operating at 10 Gbps, we achieved a maximum user-payload throughput of 1,180 megabytes per second (MBps) compared to a maximum of 680 MBps achieved with an OSA-Express3 10 GbE feature on a z196. This represents approximately a 70% increase for jumbo frames (8000 byte frames).

Measurements with mixed-direction streams workload in the same jumbo frames environment, with an OSA-Express4S 10 GbE feature, achieved a maximum user-payload throughput of 2,080 MBps with an OSA-Express4S 10 GbE feature compared to a maximum of 1,240 MBps on an OSA-Express3 10 GbE feature on z196. This represents approximately a 70% increase for jumbo frames.

Performance using standard frames : In laboratory measurements, using an OSA-Express4S 10 GbE feature in a z196 defined as CHPID type OSD with an inbound-to-the-host streams workload operating at 10 Gbps, we achieved a maximum user-payload throughput of 1,120 MBps compared to a maximum of 615 MBps achieved with an OSA-Express3 10 GbE feature on a z196. This represents approximately an 80% increase for standard frames (1492 byte frames).

Measurements with a mixed-direction streams workload in the same small frames environment, with an OSA-Express4S 10 GbE feature, achieved a maximum user-payload throughput of 1,680 MBps with an OSA-Express4S 10 GbE compared to a maximum of 1,180 MBps with an OSA-Express3 10 GbE feature on a z196. This represents approximately a 40% increase for standard frames.

OSA-Express4S performance was measured in a controlled environment using IBM Application Workload Modeler (AWM). The actual throughput or performance that any user may experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the network options and configuration, and the workload processed. One MBps represents 1,048,576 bytes per second.

Access AWM information at

<http://www.ibm.com/software/network/awm/index.html>

OSA-Express4S Gigabit Ethernet is already capable of line speed for jumbo frames and standard frames.

Parallel Sysplex® and STP enhancements

Access to a Parallel Sysplex environment

Parallel Sysplex is a synergy between hardware and software - a highly advanced technology for clustering designed to enable the aggregate capacity of multiple z/OS systems to be applied against common workloads. z/OS combined with z114, z196, z10, or z9™ machines, Coupling Facilities, Server Time Protocol (STP), and coupling links (InfiniBand, ISC-3) allows you to harness the power of multiple systems as though they were a single logical computing system.

Coupling links provide a path to transmit/receive Coupling Facility (CF) data as well as Server Time Protocol (STP) timekeeping messages. The CF data may be exchanged between z/OS and the CF or between CFs. The primary Coupling Facility configuration options are a stand-alone Coupling Facility or a CF LPAR on a server.

The primary ways to configure Coupling Facilities are:

- Stand-alone Coupling Facility, an option where the Coupling Facility LPAR running the CFCC code is not co-located with z/OS images of its Parallel Sysplex on the same System z server. This option provides for failure isolation, so that a single failure does not affect the z/OS images and the CF at the same time. The server itself may be configured with Integrated Coupling Facility specialty engines.
- Internal Coupling Facility, an option where the Coupling Facility is co-located with z/OS images of its Parallel Sysplex on the same System z server. This option can be used as a backup Coupling Facility in a data sharing environment or as the primary facility in a resource sharing environment.

InfiniBand coupling links are high-speed links, up to 6 Gbps for 12x InfiniBand and up to 5 Gbps for 1x InfiniBand:

- 12x InfiniBand for short distances - up to 150 meters (492 feet)
- 1x InfiniBand for longer distances - up to 10 km (6.2 miles) unrepeated

InterSystem Channel-3 (ISC-3) links continue to be offered, as a migration aid, supporting a link data rate of 2 Gbps. ISC-3 is supported at the same distance as 1x InfiniBand.

Internal coupling links (ICs) can also be used for internal communication between Coupling Facilities (CFs) defined in LPARs and z/OS images on the same server.

Note: The InfiniBand link data rates do not represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload. Systems Lab Services can assist your migration to PSIFB coupling links by providing services to assess the impact of the migration or to assist with the implementation of the migration.

A new generation of 12x and 1x InfiniBand coupling links

Support for demanding data-sharing workloads : Simplify your Parallel Sysplex connectivity requirements and do more with less, with a new generation of host

channel adapters (HCA3s) for coupling - the HCA3-O fanout for 12x InfiniBand (12x IFB) with improved service times using the 12x InfiniBand3 (12x IFB3) protocol and the HCA3-O LR fanout for 1x InfiniBand (1x IFB) with four ports of connectivity and optional additional subchannels for extended-distance solutions.

- 12x InfiniBand coupling links support a link data rate of 6 gigabytes per second (GBps) at distances up to 150 meters (492 feet).
- 1x IFB coupling links support a link data rate of 5 gigabits per second (Gbps) server-to-server or 2.5 or 5 Gbps when using dense wavelength division multiplexers (DWDMs). The 1x IFB coupling links support distances up to 10 km (6.2 miles) unrepeated. Greater distances are possible with an RPQ or when using qualified DWDM solutions.
- Both 12x and 1x InfiniBand coupling links also allow you to share physical links by defining multiple logical links (channel path identifiers - CHPIDs) - up to 16 CHPIDs across two ports for 12x IFB, and up to 16 CHPIDs across four ports for 1x IFB.

Note: The 12x InfiniBand HCA3-O fanout can connect to the 12x InfiniBand HCA2-O fanout on a z114, z196, or z10, but an HCA3-O fanout cannot connect to an HCA1-O fanout on a z9.

Two protocols - 12x IFB and 12x IFB3 - for the new generation of 12x InfiniBand coupling links : There are now two protocols supported by the HCA3-O for 12x IFB feature - 12x IFB and 12x IFB3.

1. **12x IFB3 protocol :** When HCA3-Os are communicating with HCA3-Os and have been defined with four or fewer CHPIDs per port, the 12x IFB3 protocol is utilized.
2. **12x IFB protocol :** If more than four CHPIDs are defined per HCA3-O port, or HCA3-O features are communicating with HCA2-O features on zEnterprise or System z10 servers, links will run with the 12x IFB protocol.

Improved service times when using the 12x IFB3 protocol : The HCA3-O feature supporting 12x InfiniBand coupling links has been designed to deliver improved service times. When no more than four CHPIDs are defined per HCA3-O port, the 12x IFB3 protocol is used. When using the 12x IFB3 protocol, synchronous service times are designed to be 40% faster than when using the 12x IFB protocol.

Improved physical connectivity with 1x InfiniBand coupling links : The HCA3-O long reach (LR) fanout for 1x InfiniBand coupling links has been designed with four ports per feature to satisfy requirements for more physical connectivity. The added connectivity will be helpful as clients migrate from InterSystem Channel-3 (ISC-3) to 1x InfiniBand coupling links.

Up to 32 subchannels (devices) per CHPID for 1x InfiniBand coupling links : To improve link utilization and coupling throughput at increased distances between the Coupling Facility (CF) and the operating system or between CFs separated by increased distances, we are now optionally supporting 32 subchannels (devices) per CHPID, versus the current 7 devices per CHPID. The additional device definition is available whether using an HCA3-O LR or HCA2-O LR feature for 1x IFB. This increase is designed to help eliminate the need to add CHPIDs or physical links in order to help improve coupling throughput at extended distance.

The HCA3-O feature (#0171) for 12x IFB and the HCA3-O LR feature (#0170) for 1x IFB are exclusive to the z196 and z114, utilize CHPID type CIB, and are supported in the z/OS environment. Refer to the [Software requirements](#) section.

Parallel Sysplex enhancements

As data sharing workloads continue to grow, the Parallel Sysplex infrastructure needs to anticipate the increased requirements for coupling resources. To do this, we have added connectivity capabilities to support larger data sharing environments.

Connectivity improvements with up to 72 coupling links: The IBM zEnterprise 114 increases the number of external coupling links allowed from 64 to 72. This allows the full configuration of 24 1x IFB links and 48 ISC-3 links to be used. In

addition, one can also configure up to 32 (internal) IC links for coupling between images defined on the same server. Having more coupling links is important to provide sufficient coupling connectivity for larger single Parallel Sysplexes, as well as for configurations where the same server hosts multiple Parallel Sysplexes and Coupling Facility images.

Connectivity improvements with 128 CHPIDs per server: To support larger Parallel Sysplexes with ever-increasing amounts of data sharing traffic to the Coupling Facility, the throughput and capacity of more coupling CHPIDs are also required. With the z114, IBM has increased the number of coupling CHPIDs per server from 64 to 128. Since IFB links allow for multiple (logical) CHPIDs over the same (physical) link, this can also allow for larger Parallel Sysplexes without requiring more coupling link hardware.

Coupling Facility Control Code Level 17

Connectivity improvements with up to 2047 structures: CFCC Level 17 increases the number of structures that can be allocated in a CFCC image from 1023 to 2047. Allowing more CF structures to be defined and used in a sysplex permits more discrete data sharing groups to operate concurrently, and can help environments requiring many structures to be defined, such as to support SAP or service providers. z114 and CFCC Level 17 provide improved serviceability of Coupling Facilities with enhanced data collection and triggering of nondisruptive CF dumps. Refer to the [Software requirements](#) section.

STP - Time synchronization for Parallel Sysplex

Server Time Protocol (STP) is designed to allow events occurring in different servers to be properly sequenced in time. STP is designed for servers that have been configured in a Parallel Sysplex or a basic sysplex (without a Coupling Facility), as well as servers that are not in a sysplex but need time synchronization.

STP is a server-wide facility that is implemented in the Licensed Internal Code (LIC), presenting a single view of time to Processor Resource/Systems Manager™ (PR/SM). STP uses a message-based protocol in which timekeeping information is passed over externally defined coupling links between servers. The coupling links that can be used to transport STP messages include 12x or 1x InfiniBand (IFB) links and InterSystem Channel-3 (ISC-3) links configured in peer mode. These can be the same links that already are being used in a Parallel Sysplex for Coupling Facility (CF) message communication. STP can also use ICB-3 and ICB-4 links on servers that support them.

The STP design introduced a concept called Coordinated Timing Network (CTN). A Coordinated Timing Network (CTN) is a collection of servers and Coupling Facilities that are time-synchronized to a time value called Coordinated Server Time.

A CTN can be configured in two ways:

- STP-only CTN, which does not require a Sysplex Timer® .
- Mixed CTN (External Time Reference (ETR) and STP), which requires a Sysplex Timer . The Sysplex Timer provides the timekeeping information in a Mixed CTN. Even though the z114 does not support attachment to a Sysplex Timer , it can participate in a Mixed CTN that has either a z10 or z9 synchronized to the Sysplex Timer . This maintains the capability for enterprises to concurrently migrate from an existing ETR network to a Mixed CTN and from a Mixed CTN to an STP-only CTN.

Server Time Protocol (STP) enhancements

Recovery enhancement : The new generation of host channel adapters (HCA3-O or HCA3-O LR), introduced for coupling, have been designed to send a reliable unambiguous "going away signal" (GOSIG) to indicate that the server on which the HCA3 is running is about to enter a failed (check stopped) state. When the GOSIG sent by the Current Time Server (CTS) in an STP-only Coordinated Timing Network (CTN) is received by the Backup Time Server (BTS), the BTS can safely take over as

the CTS without relying on the previous recovery methods of Offline Signal (OLS) in a two-server CTN or the Arbiter in a CTN with three or more servers.

This enhancement is exclusive to the z196 and z114 and is available only if you have an HCA3-O or HCA3-O LR on the CTS communicating with an HCA3-O or HCA3-O LR on the BTS. Note that the already available STP recovery design is still available for the cases when a GOSIG is not received or for other failures besides a server failure.

High-voltage DC power

In data centers today, many businesses are paying increasing electric bills and are also running out of power. High-voltage DC power, an optional feature on z114, increases the voltage directly into the system. This can save System z users 1% - 3%, on average, on their power bills without having to go through a step-down.

High-voltage DC power is added to the universal power input that is available on the z114. The voltage level is in the range of 380V - 570V DC, and requires two or four 60 A line cords, depending on configuration. These line cords (#8965, #8963) are offered on new build as well as MES orders.

Non-raised floor environment

An IBM zEnterprise 114 feature may be ordered to allow use of the z114 in a non-raised floor environment. This capability may help ease the cost of entry into the z114; a raised floor may not be necessary for some infrastructures.

The non-raised floor implementation is designed to meet all electromagnetic compatibility standards. Feature #7998 must be ordered if the z114 is to be used in a non-raised floor environment. A Bolt-down kit (#8013) is also available for use with a non-raised floor z114, providing frame stabilization and bolt-down hardware to help secure a frame to a non-raised floor.

Top exit I/O cabling and power cords

On the z114 you now have the option of ordering the infrastructure to support top exit of your fiber optic cables (ESCON , FICON , OSA, 12x InfiniBand, 1x InfiniBand, and ISC-3), your copper cables for the 1000BASE-T Ethernet features, as well as the I/O power cords.

Top exit I/O cabling and power cords are designed to provide you with additional installation options. Instead of all of your cables and power cords exiting under the server and/or under the raised floor, you now have the flexibility to choose the option that best meets the requirements of your data center. Top exit I/O cabling and power cords can also help to increase air flow. This cabling flexibility exists on both raised floor and non-raised floor installations. However for non-raised floor installations both the I/O and power must exit either out the top or out the bottom of the frame.

This option is offered on new build as well as MES orders.

Access to Capacity on Demand

Capacity on Demand - Temporary capacity

The z114 uses the Capacity on Demand (CoD) architecture implemented on the z10 . This architecture improved the capability to access and manage processing capacity on a temporary basis, providing increased flexibility for on-demand environments.

Capacity Back Up (CBU) gives you temporary access to dormant processor units (PUs), intended to replace capacity lost within the enterprise due to a disaster. CP capacity or any and all specialty engine types (ICF, IFL, SAP, zAAP, zIIP) can be added up to what the physical hardware model can contain for up to 10 days for a

test activation or 90 days for true disaster recovery. The CBU entitlement records (#6818) contain an expiration date that is established at the time of the order and is dependent upon the quantity of CBU years (#6817).

You have the capability to extend your CBU entitlements through the purchase of additional CBU years. The number of #6817 per instance of #6818 remains limited to five. Fractional years are rounded up to the near whole integer when calculating this limit. One test activation is provided for each additional CBU year added to the CBU entitlement record. The number of test activations per CBU entitlement record coincides with the number of years assigned to the CBU record. This equates to one test activation per year for each CBU entitlement purchased. Additional test activations are available in quantities of one and the number of test activations is limited at 15 per CBU record.

Capacity for Planned Events (CPE) : Temporary access to dormant processing units (PUs), intended to replace capacity lost within the enterprise due to a planned event such as a facility upgrade or system relocation. CP capacity or any and all specialty engine types (zIIP, zAAP, SAP, IFL, ICF) can be added up to what the physical hardware model can contain for up to 3 days.

On/Off Capacity on Demand (On/Off CoD) : Temporary access to dormant PUs, intended to augment the existing capacity of a given system. On/Off CoD helps you contain workload spikes that may exceed permanent capacity such that Service Level Agreements cannot be met and business conditions do not justify a permanent upgrade. As on the System z10 , On/Off CoD can be pre- or postpaid, token or nontoken.

Capacity provisioning : An installed On/Off CoD record is a necessary prerequisite for automated control of temporary capacity through z/OS MVS™ Capacity Provisioning. z/OS MVS Capacity Provisioning allows you to set up rules defining the circumstances under which additional capacity should be provisioned in order to fulfill a specific business need. The rules are based on criteria such as a specific application, the maximum additional capacity that should be activated, and time and workload conditions. This support provides a fast response to capacity changes and ensures sufficient processing power is available with the least possible delay even if workloads fluctuate. Refer to *z/OS MVS Capacity Provisioning User's Guide* (SA33-8299) for more information.

Capacity on Demand - Permanent capacity

Customer Initiated Upgrade (CIU) facility : When your business needs additional capacity quickly, Customer Initiated Upgrade (CIU) is designed to deliver it. CIU is designed to allow you to respond to sudden increased capacity requirements by requesting a z114 PU and/or memory upgrade via the web, using IBM Resource Link , and downloading and applying it to your z114 server using your system's Remote Support connection. Further, with the Express option on CIU, an upgrade may be made available for installation as fast as within a few hours after order submission.

Permanent upgrades : Orders (MESs) of all PU types and memory for System z114 servers that can be delivered by Licensed Internal Code, Configuration Control (LICCC) are eligible for CIU delivery. CIU upgrades may be performed up to the maximum available processor and memory resources on the installed server, as configured. While capacity upgrades to the server itself are concurrent, your software may not be able to take advantage of the increased capacity without performing an Initial Programming Load (IPL).

Plan ahead memory : As on the System z10 , memory upgrades can be preplanned to be nondisruptive. The preplanned memory feature adds the necessary physical memory required to support target memory sizes. If you anticipate an increase in memory requirements, a "target" logical memory size can now be specified in the configuration tool along with a "starting" logical memory size. The configuration tool then calculates the physical memory required to satisfy this target memory. Should additional physical memory be required, it is fulfilled with the preplanned memory features.

Capacity on Demand enhancements

Autorenewal of On/Off CoD allows you to choose the automatic renewal of installed On/Off CoD records to automatically update the expiration date of those records.

Order Permanent Upgrade now allows you to order additional unassigned capacity by providing you with the options to adjust the system's total and active capacity levels. You are billed for the purchase of any new capacity, however no charges are made for deactivation of the capacity as long as new capacity is purchased.

On/Off CoD Administrative Test provides you the capability to test the Capacity on Demand process, for training and API testing. This differs from the On/Off CoD no-charge test in that no capacity is activated with the test.

With your z114, for Capacity on Demand records shipped with the system, such as CBU and/or CPE, up to four come preinstalled instead of being staged on the Service Element. If more than four records are ordered with the system, none will be installed and all will be staged. IBM installation service representatives will assist you with the installation of the staged records.

HMC system support

The new functions available on the Hardware Management Console (HMC) version 2.11.1, as described, apply exclusively to z114. However, the HMC version 2.11.1 will also support the systems listed in the table below.

The 2.11.1 HMC now supports up to two 10/100/1000 Mbps Ethernet LANs.

| Family | Machine type | Firmware driver | SE version |
|---------|--------------|-----------------|------------|
| z114 | 2818 | 93 | 2.11.1 |
| z196 | 2817 | 86 | 2.11.0 |
| z10 BC | 2098 | 79 | 2.10.2 |
| z10 EC | 2097 | 79 | 2.10.2 |
| z9 BC | 2096 | 67 | 2.9.2 |
| z9 EC | 2094 | 67 | 2.9.2 |
| z890 | 2086 | 55 | 1.8.2 |
| z990 | 2084 | 55 | 1.8.2 |
| z800 | 2066 | 3G | 1.7.3 |
| z900 | 2064 | 3G | 1.7.3 |
| 9672 G6 | 9672/9674 | 26 | 1.6.2 |
| 9672 G5 | 9672/9674 | 26 | 1.6.2 |

HMC security

- Security improvements

The Monitor System Events task now allows for security logs to result in email notifications using the same type of filters and rules that are used for both Hardware and Operating System messages.

The Password Profiles task now allows for the removal of predefined password rules by the access administrator.

- Offload support for customer audit

The Audit and Log Management task adds the ability to offload a number of HMC and SE files (Audit Log, Console Event Log, Console Service History, Tasks Performed Log, and Security Logs). You can now offload to removable media as well as to remote locations via FTP. The offloaded data is available in two forms: human readable (HTML) and machine readable (XML). Offloading can be manually initiated via the new Audit and Log Management task or scheduled via the Scheduled Operations task.

- HMC user ID templates and patterns

This function provides the ability to manage adding and removing HMC users in an environment where an LDAP server is the central authority for specifying which users have access to an HMC.

- View Only User IDs/access for HMC/SE

HMC and SE User ID support allows the ability to create users who have View Only access to select tasks. The new View Only tasks are the existing tasks with minor modifications to their GUI controls to prevent any actions from being taken.

- HMC/SE Secure File Transfer Protocol (FTP) support

This function allows a secure FTP connection from an HMC/SE FTP client to a customer FTP server location.

HMC monitoring

- Environmental Efficiency Statistics task

The Environmental Efficiency Statistics task adds the ability to show historical power consumption and thermal information as well as a historical summary of processor and channel utilization. The data is presented in table form or graphical ("histogram") form, and it can also be exported to a Comma Separated Value (CSV) file.

- Monitors Dashboard task

The Monitors Dashboard task provides the following capabilities:

- A tree-based view of resources in the IBM System z
- The ability to view aggregated activity when looking at large configurations
- More detail for objects with smaller scope
- Support for new graphical ways of displaying data, such as history charts

Future changes to the existing activity task will be minimized.

HMC user interface improvements

- Classic UI and User Settings task improvements

The User Settings and Console Default User Settings console actions now have a 'Classic Style' tab. This tab allows users to change the look of the Classic UI.

- Additional control over the Toggle Lock and Details tasks

In prior releases, all Task Roles contain the Toggle Lock task and the Details tasks for the various managed objects. The Customize User Controls task now includes the Toggle Lock and Details tasks which gives the administrator the ability to remove these tasks for certain users.

- Allow setting Acceptable Status for multiple objects

In prior releases setting up Acceptable Status on a new install requires the user to set the status one object at a time. A new field, "Save as default," was added to each Acceptable Status task to allow the user to change the acceptable status for all of the currently defined objects of that type.

LPAR Controls enhancements

- Change LPAR Controls - Export to File

The Change LPAR Controls task adds the ability to export the Change LPAR Controls table data to a Comma Separated Value (CSV) file. This support is available only when a user is connected to the HMC remotely via a web browser.

- Change LPAR Controls Scheduled Operation

The Change LPAR Controls Scheduled Operation now allows the partition capping value to be specified. Details about an existing Change LPAR Controls scheduled operation can now be viewed on the Support Element.

Capacity on Demand

- Autorenewal of postpaid On/Off CoD records
Resource Link will monitor all installed On/Off CoD records. Every 90 days, Resource Link will generate a replenishment record for each installed record that will move the expiration date out 180 days. The record must be "enabled" for autorenewal. The next time the system connects to RETAIN®, a replenishment record is pushed to the system and installed. Once set, no customer renewal action is required.
- On/Off CoD Administrative Test
Resource Link reintroduces the "Order Administrative On/Off CoD test." This option supports standard order flow, including approval steps that allow for testing of processes and procedures without any resulting hardware or software charges.
- Other CoD enhancements
Ordering options were added to Resource Link to allow the purchase of unassigned CP or IFL capacity and to allow users to explore capacity needs and upgrade options prior to purchasing a permanent engine upgrade.

SNMP and CIM API enhancements

- SNMP v3.0 API support
New authentication and privacy security enhancements are added to the HMC Java™ APIs. Users need to configure SNMPv3 authentication parameters using the Customize API Settings task.
- HMC SNMP and CIM API enhancements
The SNMP and CIM APIs now allow dynamic changes to the logical partition Group Capacity setting, and allow User ID audit reports to be generated and retrieved.

HMC hardware

- Removable writable media to replace HMC DVD-RAM
A removable writable media was added as an alternate to the HMC DVD-RAM. This media is the USB Flash Memory Drive (UFD).
- HMC/SE customer and SE networks Gigabit LAN
The Ethernet switch hardware used to interconnect various parts of the system has been upgraded. The SMC Networks switch now contains 16 ports which support 1 GB speeds.

Removal of Dynamic ICF Expansion Option

- Remove support for Dynamic ICF Expansion Option
Activation Profiles removed support for Dynamic ICF expansion both across ICFs and across a pool of shared CPs. The following logical Processor Assignment selections for a Coupling Facility mode logical partition were removed:
 - Dedicated and not dedicated Internal Coupling Facility processors
 - Dedicated Internal Coupling Facility processors and not dedicated central processors

Advanced Entry Workload License Charges

Coinciding with the announcement of the IBM zEnterprise 114 (z114) server, IBM introduces a new software pricing structure called Advanced Entry Workload License Charges (AEWLC). The AEWLC pricing metric leverages the reporting mechanisms and existing Millions of Service Units per hour (MSU) tiers of the Entry Workload License Charges (EWLC) pricing metric and the Midrange Workload License Charges (MWLC) pricing metric while extending the software price / performance provided by the EWLC and MWLC tiers. AEWLC applies to eligible z/OS, z/TPF, and z/VSE operating systems and their associated middleware programs only when running on a stand-alone z114 server. For additional AEWLC information, refer to Software

Announcement [211-250](#), dated July 12, 2011 , Advanced Entry Workload License Charges offers price / performance for the IBM zEnterprise 114.

When a z114 server is in an actively coupled Parallel Sysplex , you may elect aggregated Advanced Workload License Charges (AWLC) pricing or aggregated Parallel Sysplex License Charges (PSLC) pricing, subject to all applicable terms and conditions.

For additional information on AWLC, refer to Software Announcement [210-238](#), dated July 22, 2010 , Advanced Workload License Charges offers price / performance for the IBM zEnterprise 196.

Solution Edition Series

The System z Solution Edition Series delivers mainframe qualities of service at attractive prices for new workloads. Offerings will continue to be enhanced to leverage new zEnterprise platform and IBM software functionality.

The following products or services could be purchased with this product:

Systems and Technology Group (STG) Lab Services

In support of this announcement, STG Lab Services has developed enablement services for IBM zEnterprise BladeCenter Extension (zBX). Whether the environment is z/OS , z/VM , or z/VM with Linux on System z , Lab Services can help to set up your systems and train your IT staff. Lab Services will enable the IBM zEnterprise ensemble, install and configure the necessary software components, and provide hardware connectivity services that will help to demonstrate the capabilities and functions of this new technology to help clients to quickly realize the value of the investment that you have made in the new hybrid system.

The four services available are:

- **IBM Smart Analytics Optimizer Enablement Services**
A three- to four-week service offering that provides services to assist current DB2 for z/OS clients with the planning and installation activities for the integration of the optimizer into their data center as well as configuring the hardware components necessary to enable a basic functional IBM Smart Analytics Optimizer environment and IBM System Storage unit.
- **zEnterprise Ensemble Enablement JumpStart Assistance for zBX Blades**
A three- to four-week service offering that is targeted for implementations involving IBM blades and is designed to help accelerate the implementation of an ensemble. The service will set up the z/OS and zBX infrastructure, use a test application to demonstrate the functionality of the ensemble environment, and provide education to the client with a focus on the Unified Resource Manager "Manage" functionality.
- **zEnterprise Ensemble Enablement JumpStart Assistance for z/VM**
A three- to four-week service offering that is targeted for implementations that will use the Unified Resource Manager to manage Linux virtual servers under the z/VM Hypervisor . The service will set up the z/VM and Unified Resource Manager infrastructure, use a test application to show the functionality of the ensemble environment, and provide education to the client with a focus on Unified Resource Manager "Manage" functionality.
- **zEnterprise Ensemble Enablement JumpStart Assistance for DataPower XI50z Blades**
A two-week service offering that is targeted for implementations involving the configuration and enablement of their DataPower XI50z within an IBM BladeCenter Extension (zBX) environment. Upon completion, clients will be ready to exploit the zBX technology. An education workshop is also provided that will assist them to leverage the new blades quickly and assist in the training of staff

on new functions available to the HMC to support the XI50z and zEnterprise environments.

Other z114 service offerings from STG Lab Services

STG Lab Services provides other service offerings in support of IBM zEnterprise 114 and IBM zEnterprise z196. A sampling of these services that might be applicable include:

- z/OS Infrastructure Solutions, for example
 - Parallel Sysplex
 - Hardware, Performance, and DFSMSHsm™ Health Inspections
- Middleware Solutions
 - DB2 services
 - CICS® services
 - WebSphere services
- z/OS and Linux on System z Security Services
- Linux on IBM System z Services and Training Kits
- Systems Agenda Consulting
- Cross Platform Consulting Services
 - Data Center Services
 - IBM High Availability Center of Competency briefings and education on high availability best practices
 - IBM Systems Technical Training

For more information, visit

<http://www.ibm.com/systems/services/labservices/>

Contact your IBM representative, your IBM Business Partner, or the IBM zEnterprise Enablement Services team (systemz@us.ibm.com).

IBM Implementation Services for System z

zBX and zEnterprise Unified Resource Manager services leverage skilled experts and proven methods to help you implement zBX and zEnterprise Unified Resource Manager for POWER7 blades in your zEnterprise environment. With a more reliable implementation process, you can realize faster returns on your investments and reduce the risk of business disruption. Our solution includes the enablement of an existing IBM z/OS or z/VM logical partition (LPAR) environment, a functional IBM zBX POWER7 blades test environment, and detailed planning and project management services.

For more information about the offering, refer to

<http://www.ibm.com/services/us/en/it-services/server-product-services-for-system-z.html>

IBM zEnterprise Setup and Migration Services

IBM Implementation Services for System z - zEnterprise setup and migration offerings are designed to assist you as you plan for, configure, implement, and migrate to a zEnterprise server.

The IBM zEnterprise Setup and Migration Services is comprised of a four-day standard service with three optionally selectable activities that can be acquired separately to address your specific installation requirements.

The four-day service includes:

- Onsite or remote planning session. IBM will review your overall plan, define milestones, identify potential migration scenarios, and assess your readiness for installation.
- Onsite or remote review and updating of hardware, network, and storage provisioning. IBM will:
 - Review and update your I/O configuration (IODF)
 - Assist you as needed with the OSA-Express implementation and ensure the console definitions are valid
 - Review and update your VTAM® definitions
 - Review and update your TCP/IP parameters settings
 - Review the operating system (z/OS , z/VM , z/VSE , and Linux) settings to ensure they are appropriate for the zEnterprise
 - Review the procedures for migrating the logical partition (LPAR) configurations to the new Hardware Management Console (HMC) and assist as needed
 - Identify and order the operating system (z/OS , z/VM , and z/VSE) and subsystem (DB2 , IMS™ , and CICS) maintenance
- One day onsite within 120 days of contract signing to facilitate deployment.

The following optional activities can be acquired individually, using the recommended number of the onsite daily assist feature:

- **Two weeks of services to upgrade z/OS (within coexistence policy) to the current version and release** - Applies skilled IBM resources using best practices to migrate back-level z/OS versions (within coexistence policy) to work with the zEnterprise . Contract can be expanded to include other migration activities such as subsystems, operational migration, Independent Software Vendor (ISV), user-exits migration, testing, and deployment.
- **Two weeks of services for sysplex clients exploiting Parallel Sysplex InfiniBand (PSIFB) with the zEnterprise** - Assists you in planning, implementing, and migrating your ICB and ISC-3 configurations to PSIFB.
- **Two weeks of services to assist on other installation and Implementation requirements** - Assists you on infrastructure work specific to your requirements.

For more information about the offering, contact your Business Partner or IBM representative, or send an email to

express@us.ibm.com.

Accessibility by people with disabilities

A US Section 508 Voluntary Product Accessibility Template (VPAT) containing details on accessibility compliance can be requested at

http://www.ibm.com/able/product_accessibility/index.html

Section 508 of the US Rehabilitation Act

IBM zEnterprise 114 servers are capable on delivery, when used in accordance with IBM's associated documentation, of satisfying the applicable requirements of Section 508 of the Rehabilitation Act of 1973, 29 U.S.C. Section 794d, as implemented by 36 C.F.R. Part 1194, provided that any Assistive Technology used with the product properly interoperates with it.

Product positioning

The IBM zEnterprise System offers a total systems approach to enabling fit for purpose application deployment and advanced heterogeneous systems management capabilities that are unmatched in the industry.

The announcement of the IBM zEnterprise System took workload optimized systems technology to the next level. The IBM zEnterprise System consists of the traditional subsystems - CPC, memory, I/O, and power/packaging/cooling - and new platform management firmware that is designed to enable you to integrate leading IBM server technologies to run heterogeneous workloads under the management of one system.

The IBM zEnterprise System is a powerful yet flexible server that thinks beyond technology to drive business innovation. The server is designed to offer relief from increasing IT costs, particularly those associated with server sprawl, and also provides a secure, highly virtualized platform for application development.

At the core of the zEnterprise System for both mid-sized and small enterprises is the next-generation mainframe, the zEnterprise 114 (z114). As the newest member of the System z family, the z114 brings the unique value of hybrid computing to a much broader set of businesses. The z114 is a smaller mainframe footprint than the zEnterprise 196 (z196) and continues to deliver industry-leading virtualization, security, resiliency, and connectivity technologies cost-effectively packaged as a midrange enterprise solution. Further, this new system is fully capable of participating within a z196 centric ensemble.

With extreme granularity, extensive growth options, and world-class secure data serving and transaction processing capabilities, the z114 delivers significant improvements in packaging, performance, and total system scalability over prior generations. The z114 offers unprecedented scale for small and mid-sized businesses, providing nearly 30% more capacity than its predecessor, the z10 BC.

System resources are powered by up to 14 microprocessors running at 3.8 GHz. With up to 10 customer-configurable processors and 256 GB of memory, 6.0 GB/sec InfiniBand I/O interconnect, flexible subcapacity offerings, and High Performance FICON for System z, the z114 is a great system for mission-critical work, a scalable system for hosting hundreds of virtual servers, a resilient system that can handle the 24x7 demands of today's environment, and a backup system if the need arises.

The z114 is available in two models: the M05, a single central processing drawer model, and the M10, a two-drawer model that offers additional flexibility for I/O and coupling expansion and/or increased specialty engine capability. With up to 10 configurable cores, the model naming is indicative of how many total processor units are available for customer characterization. The cores can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processors (zIIPs), Internal Coupling Facilities (ICFs), or additional System Assist Processors (SAPs), or can be used as additional spares (M10 only).

With improved processor performance, increased capacity, new hybrid computing capabilities, and significant power, space, and cooling benefits, the z114 is now a genuine data center in a box solution and a perfect fit for infrastructure simplification and true cloud computing. Unlike other proclaimed cloud solutions that are defined by a siloed architecture resource pool, the z114 leaps beyond this approach by integrating heterogeneous compute resources across System z, UNIX, and Intel that can be fully optimized and managed at the platform level to business requirements. This capability provides the unique ability to extend the reach and strategic role of existing mainframe investments across the enterprise, and an opportunity to simplify and reduce the range of skills necessary for managing the data center.

Statement of general direction

The IBM zEnterprise 196 and the zEnterprise z114 are the last System z servers to support the Power Sequence Controller (PSC) feature .

IBM intends to not offer support for the PSC (feature #6501) on future System z servers after the z196 (machine type 2817) and z114 (machine type 2818). PSC

features cannot be ordered and cannot be carried forward on upgrade to such a follow-on server.

This is a revision to the PSC statement of general direction published October 20, 2009, " IBM System z10 - Delivering Security-Rich Offerings to Protect Your Data," Hardware Announcement [109-678](#), dated October 20, 2009 . At that time we indicated the PSC feature could not be ordered on future servers. We are further clarifying that PSC features will also not be supported as carry forward on an upgrade.

The IBM zEnterprise 196 and the zEnterprise z114 will be the last servers to offer ordering of ISC-3 .

Enterprises should begin migrating from ISC-3 features (#0217, #0218, #0219) to 12x InfiniBand (#0163 - HCA2-O or #0171 - HCA3-O fanout) or 1x InfiniBand (#0168 - HCA2-O LR or #0170 - HCA3-O LR fanout) coupling links.

The IBM zEnterprise 196 and IBM zEnterprise 114 are the last System z servers to support ESCON channels.

IBM plans to not offer ESCON channels as an orderable feature on System z servers that follow the z196 (machine type 2817) and z114 (machine type 2818). In addition, ESCON channels cannot be carried forward on an upgrade to such follow-on servers. This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to feature #2323 and #2324.

System z customers should continue to eliminate ESCON channels from the mainframe wherever possible. Alternate solutions are available for connectivity to ESCON devices. IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to FICON . This offering is designed to help you to simplify and manage a single physical and operational environment - FICON channels on the mainframe with continued connectivity to ESCON devices.

The IBM zEnterprise 196 and IBM zEnterprise 114 will be the last servers to support FICON Express4 channels.

Enterprises should begin migrating from FICON Express4 channel features (#3321, #3322, #3324) to FICON Express8S channels.

The z114 and z196 will be the last servers to support OSA-Express2 features.

Enterprises should begin migrating from OSA-Express2 features (#3364, #3365, #3366) to OSA-Express3 1000BASET and OSA-Express4S features.

The z114 and z196 are planned to be the last servers to support dial-up modems for use with the Remote Support Facility (RSF), and the External Time Source (ETS) option of Server Time Protocol (STP).

The currently available Network Time Protocol (NTP) server option for ETS as well as Internet time services available using broadband connections can be used to provide the same degree of accuracy as dial-up time services. Enterprises should begin migrating from dial-up modems to Broadband for RSF connections.

Application program interfaces (APIs) to Unified Resource Manager

IBM intends to offer APIs to the Unified Resource Manager. These APIs are designed to provide access to the same underlying functions that support the Unified Resource Manager user interface and can be exploited to enable discovery, monitoring, and provisioning use cases.

IBM intends to extend the Tivoli® Integrated Service Management for System portfolio to take advantage of the zEnterprise ensemble monitoring and management capabilities provided by the Unified Resource Manager APIs.

Dynamic discovery of storage resources by the Unified Resource Manager

IBM intends to offer dynamic discovery of storage resources by Unified Resource Manager. A server administrator will be able to trigger discovery of additional storage resources through the user interface of the Unified Resource Manager.

HiperSockets integration with the IEDN

Within a zEnterprise environment, it is planned for HiperSockets to be integrated with the intraensemble data network (IEDN), extending the reach of the HiperSockets network outside of the central processor complex (CPC) to the entire ensemble, appearing as a single Layer 2 network. HiperSockets integration with the IEDN is planned to be supported in z/OS V1.13 and z/VM in a future deliverable.

HiperSockets Completion Queue

IBM plans to support transferring HiperSockets messages asynchronously, in addition to the current synchronous manner on z196 and z114. This could be especially helpful in burst situations. The Completion Queue function is designed to allow HiperSockets to transfer data synchronously if possible and asynchronously if necessary, thus combining ultra-low latency with more tolerance for traffic peaks. HiperSockets Completion Queue is planned to be supported in the z/VM and z/VSE environments.

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

Reference information

For information on POWER7 blades refer to Hardware Announcement [110-061](#), dated April 13, 2010 , IBM BladeCenter PS701 and PS702 Express Blades.

For information on IBM zEnterprise BladeCenter Extension (zBX), refer to Hardware Announcement [110-177](#), dated July 22, 2010 , System z BladeCenter Extension.

For information on IBM Smart Analytics Optimizer, refer to Hardware Announcement [210-266](#), dated July 22, 2010 , IBM Smart Analytics Optimizer for DB2 for z/OS , V1.1.

For information on firmware managing the z/VM environment refer to Software Announcement [210-234](#), dated July 22, 2010 , IBM z/VM V6.1-Virtualization Growth for zEnterprise System and Unified Resource Manager.

For information on z/OS V1.12 refer to Software Announcement [210-235](#), dated July 22, 2010 , z/OS V1.12 GA.

For information on z/OS V1.13 refer to Software Announcement [211-252](#), dated July 12, 2011 , z/OS V1.13 GA.

Product number

| Description | Machine | | |
|---------------------|---------|-------|---------|
| | Type | Model | Feature |
| IBM zEnterprise 114 | 2818 | | |

| M05 | | | |
|-----------------------------|--------------|-------|---------|
| Model M05 | | | 1135 |
| 8 GB Memory | | | 3609 |
| Description | Machine Type | Model | Feature |
| IBM zEnterprise 114 | 2818 | | |
| | | M10 | |
| Model M10 | | | 1136 |
| 152 GB Memory | | | 3624 |
| 184 GB Memory | | | 3625 |
| 216 GB Memory | | | 3626 |
| 248 GB Memory | | | 3627 |
| Description | Machine Type | Model | Feature |
| IBM zEnterprise 114 | 2818 | | |
| | | M05 | |
| | | M10 | |
| MTU 1 - D | | | 0001 |
| MTU 100 - D | | | 0002 |
| MTU 1 - V | | | 0003 |
| MTU 100 - V | | | 0004 |
| GTU 1 - D | | | 0005 |
| GTU 100 - D | | | 0006 |
| GTU 1 - V | | | 0007 |
| GTU 100 - V | | | 0008 |
| GTU 1000 - D | | | 0009 |
| GTU 1000 - V | | | 0010 |
| Manage Firmware Suite | | | 0019 |
| Automate/Adv Mgmt FW Suite | | | 0020 |
| Ensemble Membership | | | 0025 |
| Manage FW ISAOPT | | | 0039 |
| Manage FW DP | | | 0040 |
| Manage FW PWR Blade | | | 0041 |
| Manage FW System x Blade | | | 0042 |
| Automate FW ISAOPT | | | 0043 |
| Automate FW DP | | | 0044 |
| Automate FW PWR Blade | | | 0045 |
| Adv Mgmt FW System x Blade | | | 0046 |
| Automate FW IFL | | | 0052 |
| Ethernet switch | | | 0070 |
| HMC w/Dual EN | | | 0091 |
| I/O Cage ISC-D Airflow | | | 0113 |
| I/O Cage Full Card Airflow | | | 0114 |
| 1 CPE Capacity Unit | | | 0116 |
| 100 CPE Capacity Unit | | | 0117 |
| 10000 CPE Capacity Unit | | | 0118 |
| 1 CPE Capacity Unit-IFL | | | 0119 |
| 100 CPE Capacity Unit-IFL | | | 0120 |
| 1 CPE Capacity Unit-ICF | | | 0121 |
| 100 CPE Capacity Unit-ICF | | | 0122 |
| 1 CPE Capacity Unit-ZAAP | | | 0123 |
| 100 CPE Capacity Unit-ZAAP | | | 0124 |
| 1 CPE Capacity Unit-zIIP | | | 0125 |
| 100 CPE Capacity Unit-zIIP | | | 0126 |
| 1 CPE Capacity Unit-SAP | | | 0127 |
| 100 CPE Capacity Unit-SAP | | | 0128 |
| HCA2-C Fanout | | | 0162 |
| HCA2-0 fanout for 12x IFB | | | 0163 |
| Fanout Airflow | | | 0165 |
| PCIe Fanout | | | 0169 |
| HCA3-0 LR fanout for 1x IFB | | | 0170 |
| HCA3-0 fanout for 12x IFB | | | 0171 |

| | |
|-----------------------------|------|
| ISC-Mother Card | 0217 |
| ISC-Daughter Card | 0218 |
| ISC-3 link on F/C 0218 | 0219 |
| IFB-MP Daughter Card | 0326 |
| STI-A8 Mother Card | 0327 |
| PCIe Interconnect | 0400 |
| OSA-Express4S GbE LX | 0404 |
| OSA-Express4S GbE SX | 0405 |
| OSA-Express4S 10 GbE LR | 0406 |
| OSA-Express4S 10 GbE SR | 0407 |
| FICON Express8S 10km LX | 0409 |
| FICON Express8S SX | 0410 |
| Month Indicator | 0660 |
| Day Indicator | 0661 |
| Hour Indicator | 0662 |
| Minute Indicator | 0663 |
| TKE workstation | 0841 |
| Crypto Express3 | 0864 |
| TKE 7.1 LIC | 0867 |
| Crypto Express3-1P | 0871 |
| TKE addl smart cards | 0884 |
| TKE Smart Card Reader | 0885 |
| UID Label for DoD | 0998 |
| STP Enablement | 1021 |
| EMEA Special Operations | 1022 |
| 4 GB Mem DIMM (10/feature) | 1605 |
| 8 GB Mem DIMM (10/feature) | 1606 |
| 16 GB Mem DIMM (10/feature) | 1607 |
| LICCC Ship Via Net Ind | 1750 |
| 8 GB Memory Capacity Incr | 1903 |
| 8 GB Ftr Converted Memory | 1904 |
| Preplanned Memory | 1993 |
| 16-Port ESCON Card | 2323 |
| ESCON Channel port | 2324 |
| US English | 2924 |
| France | 2928 |
| German | 2929 |
| Spanish - Non Spain | 2930 |
| Spain | 2931 |
| Italian | 2932 |
| Canadian French | 2935 |
| Portuguese | 2978 |
| Brazilian Portuguese | 2979 |
| UK English | 2980 |
| Norwegian | 2983 |
| Sweden Finland | 2987 |
| Netherlands | 2988 |
| Belgian French | 2989 |
| Denmark | 2993 |
| Swiss French, German | 2997 |
| Balanced Power Plan Ahead | 3003 |
| BPR Pair Air Model | 3004 |
| BPD Pair Air Model | 3005 |
| CP-A | 3074 |
| CP-B | 3075 |
| CP-C | 3076 |
| CP-D | 3077 |
| CP-E | 3078 |
| CP-F | 3079 |
| CP-G | 3080 |
| CP-H | 3081 |
| CP-I | 3082 |
| CP-J | 3083 |
| CP-K | 3084 |
| CP-L | 3085 |
| CP-M | 3086 |

| | |
|------------------------------|------|
| CP-N | 3087 |
| CP-O | 3088 |
| CP-P | 3089 |
| CP-Q | 3090 |
| CP-R | 3091 |
| CP-S | 3092 |
| CP-T | 3093 |
| CP-U | 3094 |
| CP-V | 3095 |
| CP-W | 3096 |
| CP-X | 3097 |
| CP-Y | 3098 |
| CP-Z | 3099 |
| Internal Battery IBF-E | 3212 |
| OSA-Express3 1000BASET-EN | 3367 |
| OSA-Express3-2P 1000BASET EN | 3369 |
| IFL | 3394 |
| ICF | 3395 |
| SAP optional | 3396 |
| ZAAP | 3397 |
| ZIIP | 3398 |
| Unassigned IFL | 3399 |
| 0-Way Processor A00 | 3400 |
| 1-Way Processor A01 | 3401 |
| 1-Way Processor B01 | 3402 |
| 1-Way Processor C01 | 3403 |
| 1-Way Processor D01 | 3404 |
| 1-Way Processor E01 | 3405 |
| 1-Way Processor F01 | 3406 |
| 1-Way Processor G01 | 3407 |
| 1-Way Processor H01 | 3408 |
| 1-Way Processor I01 | 3409 |
| 1-Way Processor J01 | 3410 |
| 1-Way Processor K01 | 3411 |
| 1-Way Processor L01 | 3412 |
| 1-Way Processor M01 | 3413 |
| 1-Way Processor N01 | 3414 |
| 1-Way Processor O01 | 3415 |
| 1-Way Processor P01 | 3416 |
| 1-Way Processor Q01 | 3417 |
| 1-Way Processor R01 | 3418 |
| 1-Way Processor S01 | 3419 |
| 1-Way Processor T01 | 3420 |
| 1-Way Processor U01 | 3421 |
| 1-Way Processor V01 | 3422 |
| 1-Way Processor W01 | 3423 |
| 1-Way Processor X01 | 3424 |
| 1-Way Processor Y01 | 3425 |
| 1-Way Processor Z01 | 3426 |
| 2-Way Processor A02 | 3429 |
| 2-Way Processor B02 | 3430 |
| 2-Way Processor C02 | 3431 |
| 2-Way Processor D02 | 3432 |
| 2-Way Processor E02 | 3433 |
| 2-Way Processor F02 | 3434 |
| 2-Way Processor G02 | 3435 |
| 2-Way Processor H02 | 3436 |
| 2-Way Processor I02 | 3437 |
| 2-Way Processor J02 | 3438 |
| 2-Way Processor K02 | 3439 |
| 2-Way Processor L02 | 3440 |
| 2-Way Processor M02 | 3441 |
| 2-Way Processor N02 | 3442 |
| 2-Way Processor O02 | 3443 |
| 2-Way Processor P02 | 3444 |
| 2-Way Processor Q02 | 3445 |
| 2-Way Processor R02 | 3446 |
| 2-Way Processor S02 | 3447 |
| 2-Way Processor T02 | 3448 |

| | | |
|-----------------|-----|------|
| 2-Way Processor | U02 | 3449 |
| 2-Way Processor | V02 | 3450 |
| 2-Way Processor | W02 | 3451 |
| 2-Way Processor | X02 | 3452 |
| 2-Way Processor | Y02 | 3453 |
| 2-Way Processor | Z02 | 3454 |
| 3-Way Processor | A03 | 3457 |
| 3-Way Processor | B03 | 3458 |
| 3-Way Processor | C03 | 3459 |
| 3-Way Processor | D03 | 3460 |
| 3-Way Processor | E03 | 3461 |
| 3-Way Processor | F03 | 3462 |
| 3-Way Processor | G03 | 3463 |
| 3-Way Processor | H03 | 3464 |
| 3-Way Processor | I03 | 3465 |
| 3-Way Processor | J03 | 3466 |
| 3-Way Processor | K03 | 3467 |
| 3-Way Processor | L03 | 3468 |
| 3-Way Processor | M03 | 3469 |
| 3-Way Processor | N03 | 3470 |
| 3-Way Processor | O03 | 3471 |
| 3-Way Processor | P03 | 3472 |
| 3-Way Processor | Q03 | 3473 |
| 3-Way Processor | R03 | 3474 |
| 3-Way Processor | S03 | 3475 |
| 3-Way Processor | T03 | 3476 |
| 3-Way Processor | U03 | 3477 |
| 3-Way Processor | V03 | 3478 |
| 3-Way Processor | W03 | 3479 |
| 3-Way Processor | X03 | 3480 |
| 3-Way Processor | Y03 | 3481 |
| 3-Way Processor | Z03 | 3482 |
| 4-Way Processor | A04 | 3485 |
| 4-Way Processor | B04 | 3486 |
| 4-Way Processor | C04 | 3487 |
| 4-Way Processor | D04 | 3488 |
| 4-Way Processor | E04 | 3489 |
| 4-Way Processor | F04 | 3490 |
| 4-Way Processor | G04 | 3491 |
| 4-Way Processor | H04 | 3492 |
| 4-Way Processor | I04 | 3493 |
| 4-Way Processor | J04 | 3494 |
| 4-Way Processor | K04 | 3495 |
| 4-Way Processor | L04 | 3496 |
| 4-Way Processor | M04 | 3497 |
| 4-Way Processor | N04 | 3498 |
| 4-Way Processor | O04 | 3499 |
| 4-Way Processor | P04 | 3500 |
| 4-Way Processor | Q04 | 3501 |
| 4-Way Processor | R04 | 3502 |
| 4-Way Processor | S04 | 3503 |
| 4-Way Processor | T04 | 3504 |
| 4-Way Processor | U04 | 3505 |
| 4-Way Processor | V04 | 3506 |
| 4-Way Processor | W04 | 3507 |
| 4-Way Processor | X04 | 3508 |
| 4-Way Processor | Y04 | 3509 |
| 4-Way Processor | Z04 | 3510 |
| 5-Way Processor | A05 | 3513 |
| 5-Way Processor | B05 | 3514 |
| 5-Way Processor | C05 | 3515 |
| 5-Way Processor | D05 | 3516 |
| 5-Way Processor | E05 | 3517 |
| 5-Way Processor | F05 | 3518 |
| 5-Way Processor | G05 | 3519 |
| 5-Way Processor | H05 | 3520 |
| 5-Way Processor | I05 | 3521 |
| 5-Way Processor | J05 | 3522 |
| 5-Way Processor | K05 | 3523 |
| 5-Way Processor | L05 | 3524 |
| 5-Way Processor | M05 | 3525 |

| | | |
|----------------------------|-----------------|------|
| 5-Way Processor | N05 | 3526 |
| 5-Way Processor | O05 | 3527 |
| 5-Way Processor | P05 | 3528 |
| 5-Way Processor | Q05 | 3529 |
| 5-Way Processor | R05 | 3530 |
| 5-Way Processor | S05 | 3531 |
| 5-Way Processor | T05 | 3532 |
| 5-Way Processor | U05 | 3533 |
| 5-Way Processor | V05 | 3534 |
| 5-Way Processor | W05 | 3535 |
| 5-Way Processor | X05 | 3536 |
| 5-Way Processor | Y05 | 3537 |
| 5-Way Processor | Z05 | 3538 |
| | | |
| 16 GB | Memory | 3610 |
| 24 GB | Memory | 3611 |
| 32 GB | Memory | 3612 |
| 40 GB | Memory | 3613 |
| 48 GB | Memory | 3614 |
| 56 GB | Memory | 3615 |
| 64 GB | Memory | 3616 |
| 72 GB | Memory | 3617 |
| 80 GB | Memory | 3618 |
| 88 GB | Memory | 3619 |
| 96 GB | Memory | 3620 |
| 104 GB | Memory | 3621 |
| 112 GB | Memory | 3622 |
| 120 GB | Memory | 3623 |
| | | |
| Universal Lift Tool/Ladder | | 3759 |
| CPACF Enablement | | 3863 |
| I/O Drawer | | 4000 |
| PCIe I/O Drawer-A | | 4003 |
| | | |
| A00 | Capacity Marker | 4049 |
| A01 | Capacity Marker | 4050 |
| B01 | Capacity Marker | 4051 |
| C01 | Capacity Marker | 4052 |
| D01 | Capacity Marker | 4053 |
| E01 | Capacity Marker | 4054 |
| F01 | Capacity Marker | 4055 |
| G01 | Capacity Marker | 4056 |
| H01 | Capacity Marker | 4057 |
| I01 | Capacity Marker | 4058 |
| J01 | Capacity Marker | 4059 |
| K01 | Capacity Marker | 4060 |
| L01 | Capacity Marker | 4061 |
| M01 | Capacity Marker | 4062 |
| N01 | Capacity Marker | 4063 |
| O01 | Capacity Marker | 4064 |
| P01 | Capacity Marker | 4065 |
| Q01 | Capacity Marker | 4066 |
| R01 | Capacity Marker | 4067 |
| S01 | Capacity Marker | 4068 |
| T01 | Capacity Marker | 4069 |
| U01 | Capacity Marker | 4070 |
| V01 | Capacity Marker | 4071 |
| W01 | Capacity Marker | 4072 |
| X01 | Capacity Marker | 4073 |
| Y01 | Capacity Marker | 4074 |
| Z01 | Capacity Marker | 4075 |
| | | |
| A02 | Capacity Marker | 4078 |
| B02 | Capacity Marker | 4079 |
| C02 | Capacity Marker | 4080 |
| D02 | Capacity Marker | 4081 |
| E02 | Capacity Marker | 4082 |
| F02 | Capacity Marker | 4083 |
| G02 | Capacity Marker | 4084 |
| H02 | Capacity Marker | 4085 |
| I02 | Capacity Marker | 4086 |
| J02 | Capacity Marker | 4087 |
| K02 | Capacity Marker | 4088 |
| L02 | Capacity Marker | 4089 |

| | |
|---------------------|------|
| M02 Capacity Marker | 4090 |
| N02 Capacity Marker | 4091 |
| O02 Capacity Marker | 4092 |
| P02 Capacity Marker | 4093 |
| Q02 Capacity Marker | 4094 |
| R02 Capacity Marker | 4095 |
| S02 Capacity Marker | 4096 |
| T02 Capacity Marker | 4097 |
| U02 Capacity Marker | 4098 |
| V02 Capacity Marker | 4099 |
| W02 Capacity Marker | 4100 |
| X02 Capacity Marker | 4101 |
| Y02 Capacity Marker | 4102 |
| Z02 Capacity Marker | 4103 |
| | |
| A03 Capacity Marker | 4106 |
| B03 Capacity Marker | 4107 |
| C03 Capacity Marker | 4108 |
| D03 Capacity Marker | 4109 |
| E03 Capacity Marker | 4110 |
| F03 Capacity Marker | 4111 |
| G03 Capacity Marker | 4112 |
| H03 Capacity Marker | 4113 |
| I03 Capacity Marker | 4114 |
| J03 Capacity Marker | 4115 |
| K03 Capacity Marker | 4116 |
| L03 Capacity Marker | 4117 |
| M03 Capacity Marker | 4118 |
| N03 Capacity Marker | 4119 |
| O03 Capacity Marker | 4120 |
| P03 Capacity Marker | 4121 |
| Q03 Capacity Marker | 4122 |
| R03 Capacity Marker | 4123 |
| S03 Capacity Marker | 4124 |
| T03 Capacity Marker | 4125 |
| U03 Capacity Marker | 4126 |
| V03 Capacity Marker | 4127 |
| W03 Capacity Marker | 4128 |
| X03 Capacity Marker | 4129 |
| Y03 Capacity Marker | 4130 |
| Z03 Capacity Marker | 4131 |
| | |
| A04 Capacity Marker | 4134 |
| B04 Capacity Marker | 4135 |
| C04 Capacity Marker | 4136 |
| D04 Capacity Marker | 4137 |
| E04 Capacity Marker | 4138 |
| F04 Capacity Marker | 4139 |
| G04 Capacity Marker | 4140 |
| H04 Capacity Marker | 4141 |
| I04 Capacity Marker | 4142 |
| J04 Capacity Marker | 4143 |
| K04 Capacity Marker | 4144 |
| L04 Capacity Marker | 4145 |
| M04 Capacity Marker | 4146 |
| N04 Capacity Marker | 4147 |
| O04 Capacity Marker | 4148 |
| P04 Capacity Marker | 4149 |
| Q04 Capacity Marker | 4150 |
| R04 Capacity Marker | 4151 |
| S04 Capacity Marker | 4152 |
| T04 Capacity Marker | 4153 |
| U04 Capacity Marker | 4154 |
| V04 Capacity Marker | 4155 |
| W04 Capacity Marker | 4156 |
| X04 Capacity Marker | 4157 |
| Y04 Capacity Marker | 4158 |
| Z04 Capacity Marker | 4159 |
| | |
| A05 Capacity Marker | 4162 |
| B05 Capacity Marker | 4163 |
| C05 Capacity Marker | 4164 |
| D05 Capacity Marker | 4165 |
| E05 Capacity Marker | 4166 |

| | |
|-----------------------------|------|
| F05 Capacity Marker | 4167 |
| G05 Capacity Marker | 4168 |
| H05 Capacity Marker | 4169 |
| I05 Capacity Marker | 4170 |
| J05 Capacity Marker | 4171 |
| K05 Capacity Marker | 4172 |
| L05 Capacity Marker | 4173 |
| M05 Capacity Marker | 4174 |
| N05 Capacity Marker | 4175 |
| O05 Capacity Marker | 4176 |
| P05 Capacity Marker | 4177 |
| Q05 Capacity Marker | 4178 |
| R05 Capacity Marker | 4179 |
| S05 Capacity Marker | 4180 |
| T05 Capacity Marker | 4181 |
| U05 Capacity Marker | 4182 |
| V05 Capacity Marker | 4183 |
| W05 Capacity Marker | 4184 |
| X05 Capacity Marker | 4185 |
| Y05 Capacity Marker | 4186 |
| Z05 Capacity Marker | 4187 |
| | |
| Luxembourg-Belgium ordered | 5560 |
| Iceland-Ordered in Denmark | 5561 |
| China-Ordered in Hong Kong | 5562 |
| Flat panel display | 6096 |
| Power Sequence Controller | 6501 |
| | |
| Additional CBU Test | 6805 |
| Total CBU Years Ordered | 6817 |
| CBU Records Ordered | 6818 |
| Single CBU CP Year | 6820 |
| 25 CBU CP Year | 6821 |
| Single CBU IFL Year | 6822 |
| 25 CBU IFL Year | 6823 |
| Single CBU ICF Year | 6824 |
| 25 CBU ICF Year | 6825 |
| Single CBU ZAAP Year | 6826 |
| 25 CBU ZAAP Year | 6827 |
| Single CBU ZIIP Year | 6828 |
| 25 CBU ZIIP Year | 6829 |
| Single CBU SAP Year | 6830 |
| 25 CBU SAP Year | 6831 |
| CBU Replenishment | 6832 |
| Capacity for Planned Event | 6833 |
| OPO Sales Flag | 6835 |
| OPO Sales Flag Alteration | 6836 |
| Feature Converted CBU CP | 6837 |
| Feature Converted CBU IFL | 6838 |
| Feature Converted CBU ICF | 6839 |
| Feature Converted CBU ZAAP | 6840 |
| Feature Converted CBU ZIIP | 6841 |
| Feature Converted CBU SAP | 6842 |
| | |
| MT-RJ 8.5 ft harness | 7830 |
| FQC Bracket & Mounting Hdw | 7832 |
| MT-RJ 3.5 ft harness | 7833 |
| MT-RJ 5 ft harness | 7834 |
| MT-RJ 6 ft harness | 7835 |
| LC DUP 6.6 ft harness | 7836 |
| LC DUP 8.5ft harness | 7837 |
| Top Exit w/Power | 7901 |
| Top Exit Cabling | 7920 |
| Side Covers | 7921 |
| Non Raised Floor Support | 7998 |
| 3n1 Bolt Down Kit | 8012 |
| Bolt-Down Kit, Non Raised F | 8013 |
| | |
| 14ft 380-520V DC Cut Cord | 8964 |
| 14ft HiLoV 3Ph Cut Top Exit | 8970 |
| 14ft HiLoV 1Ph Cut Top Exit | 8972 |
| 14ft 380-520V DC Cut Top Ex | 8974 |
| 14ft Hi/LoV 3 Ph Cut Cord | 8988 |
| 14ft HiLoV 1 Ph Cut Cord | 8991 |

| | |
|------------------------------|------|
| 14ft Hiv 3Ph Cut line-LSZH | 8998 |
| 14ft 380-520V DC Line Cord | 8966 |
| 6ft 200V 1Ph Cord Top Exit | 8971 |
| 6ft 380-520VDC Top Exit | 8973 |
| 6ft 200V 3Ph Cord Top Exit | 8975 |
| 14ft 200V 30A 3 Ph Cord | 8987 |
| 14ft 200V 30A 1 Ph Line Cord | 8990 |
| 6ft 480V 3Ph Cord Top Exit | 8969 |
| 14ft 480V 30A 3PH Cord | 8983 |
| Multi Order Ship Flag | 9000 |
| Multi Order Rec Only Flag-NB | 9001 |
| Multi Order Rec Only MES | 9002 |
| RPO Action Flag | 9003 |
| Downgraded PUS Per Request | 9004 |
| Site Prep | 9879 |
| On/Off CoD Act IFL Day | 9888 |
| On/Off CoD Act ICF Day | 9889 |
| On/Off CoD Act ZAAP Day | 9893 |
| On/Off CoD authorization | 9896 |
| On/Off CoD Act Cap CP Day | 9897 |
| Perm upgr authorization | 9898 |
| CIU Activation (Flag) | 9899 |
| On Line CoD Buying (Flag) | 9900 |
| On/Off CoD Act zIIP Day | 9908 |
| On/Off CoD Act SAP Day | 9909 |
| CBU authorization | 9910 |
| CPE authorization | 9912 |
| OPO Sales authorization | 9913 |
| 1 MSU day | 9917 |
| 100 MSU days | 9918 |
| 10000 MSU days | 9919 |
| 1 IFL day | 9920 |
| 100 IFL days | 9921 |
| 1 ICF day | 9922 |
| 100 ICF days | 9923 |
| 1 zIIP day | 9924 |
| 100 zIIP days | 9925 |
| 1 ZAAP day | 9926 |
| 100 ZAAP days | 9927 |
| 1 SAP day | 9928 |
| 100 SAP days | 9929 |
| Site Tool Kit | 9969 |
| Height Reduce Ship | 9975 |
| Height Reduce for Return | 9976 |

The following features are not orderable on the IBM zEnterprise 114 models. If they are installed at the time of an upgrade to the IBM zEnterprise 114 they may be retained.

| | |
|---------------------------|------|
| HMC w/Dual EN | 0084 |
| SE-EN Switch (former HUB) | 0089 |
| HMC w/Dual EN | 0090 |
| HCA2-O LR fanout | 0168 |
| FICON Express4 - 2C SX | 3318 |
| FICON Express4 10KM LX | 3321 |
| FICON Express4 SX | 3322 |
| FICON Express4-2C 4KM LX | 3323 |
| FICON Express4 4KM LX | 3324 |

| | |
|---------------------------|------|
| FICON Express8 10KM LX | 3325 |
| FICON Express8 SX | 3326 |
| OSA-Express3 GbE LX | 3362 |
| OSA-Express3 GbE SX | 3363 |
| OSA-Express2 GbE LX | 3364 |
| OSA-Express2 GbE SX | 3365 |
| OSA-Express2 1000BASET-EN | 3366 |
| OSA-Express3 10 GbE LR | 3370 |
| OSA-Express3 10 GbE SR | 3371 |
| OSA-Express3-2P GbE SX | 3373 |
| 17 inch panel display | 6094 |
| 21 inch panel display | 6095 |
| 6ft 480V 30A 3PH Cord | 8984 |
| 6ft 200V 30A 3PH Cord | 8986 |
| 6ft 200V 30A 1PH Cord | 8989 |
| 14ft HiV 1 Ph Cut-LSZH | 8999 |

Notes:

1. Memory DIMMs do NOT carry forward.
2. Support Elements do NOT carry forward.
3. FICON Express , feature numbers 2319 and 2320, are NOT supported on the z114.
4. FICON Express2, feature numbers 3319 and 3320, are NOT supported on the z114.
5. Crypto Express2, feature number 0863, is NOT supported on the z114.
6. OSA-Express2 10 GbE LR, feature number 3368, is NOT supported on z114.

Model conversions

Model Conversions - Hardware upgrades

| From M/T | Model | To M/T | Model | Description |
|-------------|-------|-----------|---------|----------------------|
| 2096 | R07 | 2818 | M05 (*) | 2096 R07 to 2818 M05 |
| 2096 | R07 | 2818 | M10 (*) | 2096 R07 to 2818 M10 |
| 2096 | S07 | 2818 | M05 (*) | 2096 S07 to 2818 M05 |
| 2096 | S07 | 2818 | M10 (*) | 2096 S07 to 2818 M10 |
| 2098 | E10 | 2818 | M05 (*) | 2098 E10 to 2818 M05 |
| 2098 | E10 | 2818 | M10 (*) | 2098 E10 to 2818 M10 |
| 2818 | M05 | 2818 | M10 (*) | 2818 M05 to 2818 M10 |
| 2818 | M10 | 2817 | M15 (*) | 2818 M10 to 2817 M15 |

(*) Parts removed or replaced become the property of IBM and must be returned.

Feature conversions

The feature conversion list for IBM zEnterprise 114 is available now in the *Library* section of Resource Link . This list can be obtained at Resource Link by accessing the following website

<http://www.ibm.com/servers/resourcelink>

Using the instructions on the Resource Link panels, obtain a user ID and password. Resource Link has been designed for easy access and navigation.

Business Partner information

If you are a Direct Reseller - System Reseller acquiring products from IBM, you may link directly to Business Partner information for this announcement. A PartnerWorld ID and password are required (use IBM ID).

<https://www.ibm.com/partnerworld/mem/sla.jsp?num=111-136>

Publications

The following publications are available now in the *Library* section of Resource Link :

| Title | Order number |
|--|--------------|
| Functional Matrix | ZSW0-1335 |
| IOCP User's Guide | SB10-7037 |
| PR/SM Planning Guide | SB10-7155 |
| CHPID Mapping Tool User's Guide | GC28-6900 |
| z114 Installation Manual for Physical Planning | GC28-6907 |
| z114 System Overview | SA22-1087 |

The following publications are shipped with the product and are available in the *Library* section of Resource Link :

| Title | Order number |
|--|--------------|
| Systems Safety Notices | G229-9054 |
| License Agreement for Machine Code Addendum for Elliptical Curve Cryptography | GC27-2612 |
| System z Statement of Limited warranty | GC28-6883 |
| Service Guide for TKE Workstations (Version 7.0) | GC28-6901 |
| z114 Installation Manual | GC28-6902 |
| z114 Service Guide | GC28-6903 |
| z114 Safety Inspection | GC28-6904 |
| License Agreement for Machine Code | SC28-6872 |
| Systems Environmental Notices and User Guide | Z125-5823 |

The following publications will be available at planned availability in the *Library* section of Resource Link :

| Title | Order number |
|--|--------------|
| System z Application Programming Interfaces for Java | API-JAVA |
| Planning for Fiber Optic Links | GA23-0367 |
| Ensemble Performance Management Guide | GC27-2607 |
| Ensemble Planning and Configuring Guide | GC27-2608 |
| Introduction to Ensembles | GC27-2609 |
| Service Guide for HMCs and SES | GC28-6861 |
| z114 Parts Catalog | GC28-6908 |
| OSA-Express Customer Guide and Reference | SA22-7935 |
| ESCON I/O Interface Physical Layer | SA23-0394 |
| Coupling Links I/O Interface Physical Layer | SA23-0395 |
| Fibre Channel Connection (FICON) I/O Physical Layer | SA24-7172 |
| System z Application Programming Interfaces | SB10-7030 |
| System z CIM Management Interface | SB10-7154 |
| Stand-Alone IOCP User's Guide | SB10-7152 |
| ESCON and FICON CTC Reference | SB10-7034 |
| HMC Operations Guide for Ensembles (v2.11.1) | SC27-2615 |
| Capacity On Demand User's Guide | SC28-2605 |
| SCSI IPL - Machine Loader Messages | SC28-6839 |
| Hardware Management Console Operations Guide (v2.11.1) | SC28-6905 |
| Support Element Operations Guide (v2.11.1) | SC28-6906 |
| Maintenance Information for Fiber Optic Links | SY27-2597 |
| Fiber Optic Cleaning Procedures | SY27-2604 |

Publications can be obtained at Resource Link by accessing the following website

<http://www.ibm.com/servers/resourcelink>

Using the instructions on the Resource Link panels, obtain a user ID and password. Resource Link has been designed for easy access and navigation.

The following Redbook publication is available now:

| Title | Order number |
|-------|--------------|
|-------|--------------|

The following Redbook publications will provide additional information, once they become available:

| Title | Order number |
|---|--------------|
| IBM zEnterprise 114 Technical Guide | SG24-7833 |
| System z Connectivity Handbook | SG24-5444 |
| IBM zEnterprise 114 Configuration Setup | SG24-7834 |
| zEnterprise Platform Management | SG24-7835 |

For other IBM Redbooks® publications, refer to

<http://www.redbooks.ibm.com/>

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

<http://publib14.boulder.ibm.com/infocenter/systems>

IBM Publications Center Portal

<http://www.ibm.com/shop/publications/order>

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, as well as payment options via credit card. A large number of publications are available online in various file formats, which can currently be downloaded free of charge.

Services

Global Technology Services

IBM services include business consulting, outsourcing, hosting services, applications, and other technology management.

These services help you learn about, plan, install, manage, or optimize your IT infrastructure to be an on-demand business. They can help you integrate your high-speed networks, storage systems, application servers, wireless protocols, and an array of platforms, middleware, and communications software for IBM and many non-IBM offerings. IBM is your one-stop shop for IT support needs.

For details on available services, contact your IBM representative or visit

<http://www.ibm.com/services/>

For details on available IBM Business Continuity and Recovery Services, contact your IBM representative or visit

<http://www.ibm.com/services/continuity>

For details on education offerings related to specific products, visit

<http://www.ibm.com/services/learning/index.html>

Select your country, and then select the product as the category.

Technical information

Specified operating environment

Physical specifications

Dimensions (rounded to the nearest 0.1 inch or 0.1 centimeter):

| | Depth | Width | Height |
|----------------------------------|-------|-------|--------|
| Without top exit I/O | | | |
| - Inches | 62.0 | 30.9 | 79.3 |
| - Centimeter | 157.5 | 78.4 | 201.3 |
| With top exit I/O | | | |
| - Inches (O/H IO cable exit) | 62.0 | 36.9 | 79.3 |
| - Centimeter (O/H IO cable exit) | 157.5 | 93.6 | 201.3 |

Approximate weight:

| | Maximum System Model M05 | Maximum System Model M10 |
|----------------------------|--------------------------------|--------------------------------|
| System without IBF Feature | | |
| - kg | 813.6 | 859.1 |
| - lb | 1790.0 | 1890.0 |
| System with IBF Feature | | |
| - kg | 909.0 | 954.5 |
| - lb | 2000.0 | 2100.0 |

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

The DC Power feature has no effect on the machine dimensions and weight.

Operating environment

- Operating temperature and altitude:
 - Minimum temperature = 10°C (50°F)
 - Maximum temperature = 35°C (95°F) up to 3000 feet, then a derating of 1°C (1.8°F) for every 1000 feet above 3000 feet
 - Maximum altitude: 10,000 feet (28°C (82.4°C) with derating)
- Relative humidity: 20% to 80% RH, maximum dewpoint = 21°C (69.8°F)
- Electric power (absolute maximum)
 - M05: 5.5 kW
 - M10: 7.5 kW

Acoustical category 1B

- M05: 5.2 bels
- M10: 7.4 bels

Capacity of exhaust

- 2100 cubic meters/hr (1235 CFM) normal room
- 2900 cubic meters/hr (1700 CFM)

Ambient > 28°C and or altitude > 3000 feet

Acoustical noise level (nominal conditions):

- Maximum Configuration: Model M10
 - Declared A-Weighted Sound Power Level, LWAd(B)
 - Declared A-Weighted Sound Pressure Level, LpAm(dB)

Leakage and starting current: 60 mA / 170A (less than on millisecond)

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

Hardware requirements

The hardware requirements for the machine type, models, features, and functions contained in this announcement are identified.

Machine Change Levels (MCLs) are required.

Descriptions of the MCLs are available now on Resource Link .

Access Resource Link at

<http://www.ibm.com/servers/resourcelink>

MCLs are designed to be applied concurrently. Contact IBM service personnel for further information.

Required Driver Level and MCLs to be made available at planned availability:

HMC (V2.11.1) plus MCLs

Support Element (V2.11.1) plus MCLs

You should review the PSP buckets for minimum Machine Change Levels (MCLs) and software PTF levels before IPLing operating systems.

Peripheral hardware and device attachments

IBM devices previously attached to IBM System z9 , IBM System z10 , and zEnterprise servers are supported for attachment to IBM zEnterprise 114 channels, unless otherwise noted. The subject I/O devices must meet ESCON or FICON architecture requirements to be supported. I/O devices that meet OEMI architecture requirements are supported only using an external converter. Prerequisite Engineering Change Levels may be required. For further detail, contact IBM service personnel.

While the z114 supports devices as described above, IBM does not commit to provide support or service for an IBM device that has reached its End of Service effective date as announced by IBM .

Note: IBM cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions regarding the capabilities of non-IBM products should be addressed to the suppliers of those products.

Information on switches and directors qualified for IBM System z FICON and FCP channels can be found in the *Library* section of Resource Link .

<http://www.ibm.com/servers/resourcelink/>

Software requirements

Operating system support

Listed are the operating systems and the minimum versions and releases supported by IBM zEnterprise 114, its functions, and its features. Select the releases appropriate to your operating system environments.

Note: Refer to the z/OS , z/VM , z/VSE subsets of the 2818DEVICE Preventive Service Planning (PSP) bucket prior to installing a z114.

IBM zEnterprise 114 requires at a minimum:

- z/OS V1.13 with PTFs.

- z/OS V1.12 with PTFs.
- z/OS V1.11 with PTFs.
- z/OS V1.10 (1) with PTFs.
- z/OS V1.9 (2) with the IBM Lifecycle Extension for z/OS V1.9 (5646-A01) with required maintenance.
- z/OS V1.8 (3) with the IBM Lifecycle Extension for z/OS V1.8 (5638-A01) with required maintenance.
- z/VM V6.1 with PTFs.
- z/VM V5.4 with PTFs.
- z/VSE V4.3 with PTFs.
- z/VSE V4.2 with PTFs.
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL6.

CP Assist for Cryptographic Function (CPACF) (#3863) enhancements require at a minimum:

- z/OS V1.10, z/OS V1.11, or z/OS V1.12 with the Cryptographic Support for z/OS V1R10-V1R12 web deliverable. This may be obtained at <http://www.ibm.com/systems/z/os/zos/downloads/>
- z/VM V5.4 with PTFs for guest exploitation.
- z/VSE V4.2 and later supports the CPACF features with the functionality supported on IBM System z10 .
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 11 SP1.
 - Red Hat RHEL 6.1.
 - For Message-Security-Assist-Extension 4 exploitation, IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

Crypto Express3 requires at a minimum:

- z/OS V1.10, z/OS V1.11, or z/OS V1.12 with the Cryptographic Support for z/OS V1R10-V1R12 web deliverable. This may be obtained at <http://www.ibm.com/systems/z/os/zos/downloads/>
- z/VM V5.4 with PTFs for guest exploitation.
 - Note: z/VM V5.4 and later support clear-and secure-key operations.
- z/VSE V4.2 and IBM TCP/IP for VSE/ESA™ V1.5.0 with PTFs.
 - Note: z/VSE supports clear-key RSA operations only.
- z/TPF V1.1 (acceleration mode only).
- Linux on System z distributions:
 - For toleration:
 - Novell SUSE SLES10 SP3 and SLES 11.
 - Red Hat RHEL 5.4 and RHEL 6.0.
 - For exploitation:
 - Novell SUSE SLES11 SP1.
 - Red Hat RHEL 6.1.

Crypto Express3 (when defined as a coprocessor) supporting **ANSI X9.8 PIN security, Enhanced CCA key wrapping, Secure Keyed-Hash Message**

Authentication Code (HMAC), and initial Elliptic Curve Cryptography (ECC) functions requires one of the following:

- z/OS V1.11, z/OS V1.12, or z/OS V1.13 with the Cryptographic Support for z/OS V1R11-V1R13 web deliverable (planned to be available September 9, 2011).
- z/OS V1.10, z/OS V1.11, or z/OS V1.12 with the Cryptographic Support for z/OS V1R10-V1R12 web deliverable with PTFs. The Cryptographic web deliverables are available at
<http://www.ibm.com/systems/z/os/zos/downloads/>
- z/OS V1.13 with PTFs.
- z/VM 5.4 with PTFs for guest exploitation.

Crypto Express3 (when defined as a coprocessor) supporting **PKA RSA OAEP with SHA-256 algorithm** requires one of the following:

- z/OS V1.10 (1), z/OS V1.11, or z/OS V1.12 with the Cryptographic Support for z/OS V1R10-V1R12 web deliverable with a PTF (planned to be available September 9, 2011).
- z/OS V1.11, z/OS V1.12, or z/OS V1.13 with the Cryptographic Support for z/OS V1R11-V1R13 web deliverable (planned to be available September 9, 2011).
- z/OS V1.13 with a PTF (planned to be available September 9, 2011).
- z/VM 5.4 with PTFs for guest exploitation.

Crypto Express3 (when defined as a coprocessor) with Expanded support for AES algorithm, Enhanced ANSI TR-31 Secure Key Exchange, PIN block decimalization table protection, and Additional Elliptic Curve Cryptography (ECC) functions requires at a minimum:

- z/OS V1.11, z/OS V1.12, or z/OS V1.13 with the Cryptographic Support for z/OS V1R11-V1R13 web deliverable (planned to be available September 9, 2011). The Cryptographic web deliverables are available at
<http://www.ibm.com/systems/z/os/zos/downloads/>
- z/VM V5.4 with PTFs for guest exploitation.
- Linux on System z distributions:
 - Support is planned for a future release of Common Cryptographic Architecture Support Program for Linux on System z .

Crypto Express3 (when defined as an accelerator) supporting **Modulus Exponent (ME) and Chinese Remainder Theorem (CRT)** requires at a minimum:

- z/OS V1.11, z/OS V1.12, or z/OS V1.13 with the Cryptographic Support for z/OS V1R11 - V1R13 web deliverable (planned to be available September 9, 2011). The Cryptographic web deliverables may be obtained at
<http://www.ibm.com/systems/z/os/zos/downloads/>
- z/VM V5.4 with PTFs for guest exploitation.
- Linux on System z distributions:
 - For RSA encryption and decryption with key lengths up to 2048-bits:
 - Novell SUSE SLES10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.
 - For RSA encryption and decryption with key lengths up to 4096-bits:
 - IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

IPL from an alternate subchannel set requires at a minimum:

- z/OS V1.11 or V1.12 with PTFs.
- z/OS V1.13.

FICON Express8S (CHPID type FC) when utilizing native FICON or Channel-To-Channel (CTC), requires at a minimum:

- z/OS V1.8 (3).
- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.

FICON Express8S (CHPID type FC) for support of zHPF single-track operations requires at a minimum:

- z/OS V1.11.
- z/OS V1.10 (1) with PTFs.
- z/OS V1.9 (2) with PTFs.
- z/OS V1.8 (3) with PTFs.
- Linux on System z distributions:
 - Novell SUSE SLES 11 SP1.
 - Red Hat RHEL 6.

FICON Express8S (CHPID type FC) for support of zHPF multitrack operations requires at a minimum:

- z/OS V1.11 with PTFs.
- z/OS V1.10 (1) with PTFs.
- z/OS V1.9 (2) with PTFs.
- Red Hat RHEL 6.1
- Linux on System z distributions:

IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

FICON Express8S (CHPID type FCP) for support of SCSI devices requires at a minimum:

- z/VM V5.4.
- z/VSE V4.2.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.

FICON Express8S (CHPID type FCP) support of hardware data router requires at a minimum:

- Linux on System z distributions:
 - IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

T10-DIF support by the FICON Express8S and FICON Express8 features when defined as CHPID type FCP requires at a minimum:

- z/VM 5.4 with PTFs for guest exploitation
- Linux on System z distributions:

IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

OSA-Express4S GbE LX (#0404) and GbE SX (#0405) require at minimum:

CHPID type OSD with exploitation of two ports per CHPID:

- z/OS V1.10 (1).
- z/OS V1.9 (2) with PTFs.
- z/OS V1.8 (3) with PTFs.
- z/VM V6.1.
- z/VM V5.4 with PTFs.
- z/VSE V4.2.
- z/TPF V1.1 PUT 4 with PTFs.
- Linux on System z distributions:
 - Novell SUSE SLES 10 SP2 and SLES 11.
 - Red Hat RHEL 5.2 and RHEL 6.0.

CHPID types OSD without maximum port exploitation (one port on the PCIe adapter is available for use):

- z/OS V1.8 (3).
- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.

OSA-Express4S 10 GbE LR (#0406) and 10 GbE SR (#0407) require at a minimum:

CHPID type OSD:

- z/OS V1.8 (3).
- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.

CHPID type OSX for access control to the intraensemble data network (IEDN) from z114 to Unified Resource Manager functions:

- z/OS V1.12 with PTFs.
- z/OS V1.11 with PTFs.
- z/OS V1.10 (1) with PTFs.
- z/VM V6.1 with PTFs.
- z/VM V5.4 with PTFs to define, modify, and delete OSX CHPID types when z/VM is the controlling LPAR for dynamic I/O.
- z/VSE V5.1 (when available).
- z/TPF V1.1 PUT 4 with PTFs.
- Linux on System z :

- Novell SLES 10 SP4 and SLES 11 SP1 (maintenance update).
- Red Hat RHEL 5.6 and RHEL 6.0.

OSA-Express4S Large Send for IPv6 packets (CHPID types OSD and OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 for guest exploitation.

OSA-Express4S checksum offload for IPv6 packets (CHPID types OSD and OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 for guest exploitation.

OSA-Express4S checksum offload for LPAR-to-LPAR traffic for IPv4 and IPv6 packets (CHPID types OSD and OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 for guest exploitation.

Inbound workload queuing for Enterprise Extender for the OSA-Express4S and OSA-Express3 features (CHPID types OSD or OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 with PTFs for guest exploitation.

IBM zEnterprise Unified Resource Manager requires at a minimum:

- z/OS V1.13 with PTFs.
- z/OS V1.12 with PTFs.
- z/OS V1.11 with PTFs.
- z/OS V1.10(1) with PTFs.

Older versions of z/OS will tolerate, but do not exploit, the IBM zEnterprise Unified Resource Manager.

- z/VM 6.1 with PTFs for enablement of virtual server life cycle management and support for managing real and virtual networking resources by the Unified Resource Manager.

z/OS discovery and autoconfiguration (zDAC) for FICON Express8 and FICON Express4 channels (CHPID type FC) requires at a minimum:

- z/OS V1.12 with PTFs.

Up to 32 HiperSockets (CHPID type IQD) requires at a minimum:

- z/OS V1.12 with PTFs.
- z/OS V1.11 with PTFs.
- z/OS V1.10 (1) with PTFs.
- z/VM V5.4 with PTFs.
- z/VSE V4.2.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.

HiperSockets network traffic analyzer (HS NTA) on IBM zEnterprise 114 and System z10 requires at a minimum:

- z/VM V5.4 with PTFs for guest exploitation.
- Linux on System z distributions:
 - Novell SUSE SLES 11 SP1.
 - Red Hat RHEL 6.0.

Note: OSA-Express3 GbE and 1000BASE-T contain two ports per PCIe adapter. The ports "share" one channel path identifier (CHPID). There may be two or four ports per feature, depending upon feature number. OSA-Express3 10 GbE SR and LR contain one port per PCIe adapter. There are two ports per feature.

| CHPID type | Applicable features | Purpose/Traffic |
|------------|---------------------|---|
| OSC | 1000BASE-T | OSA-Integrated Console Controller (OSA-ICC) Supports TN3270E, non-SNA DFT to IPL CPCs & LPS. |
| OSD | All | Queue Direct Input/Output (QDIO) architecture; TCP/IP traffic when Layer 3 (uses IP address) and Protocol-independent when Layer 2 (uses MAC address). |
| OSE | 1000BASE-T | Non-QDIO; For SNA/APPN/HPR traffic and TCP/IP passthru traffic. |
| OSM | 1000BASE-T | OSA-Express for Unified Resource Manager. Connectivity to intranode management network (INMN) from z114 to Unified Resource Manager functions. |
| OSN | GbE 1000BASE-T | OSA-Express for NCP; Appears to OS as a device supporting channel data link control (CDLC) protocol. Enables Network Control Program (NCP) channel-related functions such as loading and dumping to NCP. Provides LP-to-LP connectivity OS to IBM Communication Controller for Linux (CCL). Note: OSN CHPID type is not supported on OSA-Express4S GbE features. |
| OSX | 10 GbE | OSA-Express for zBX. Provides connectivity and access control to the intraensemble data network (IEDN) from z114 to Unified Resource Manager functions. |

OSA-Express3 1000BASE-T (#3367 and #3369) requires at minimum:

CHPID type OSC supporting TN3270E and non-SNA DFT:

- z/OS V1.8 (3).
- z/VM V5.4.
- z/VSE V4.2.

CHPID type OSD with exploitation of all four ports:

- z/OS V1.10 (1).
- z/OS V1.9 (2) with PTFs.
- z/OS V1.8 (3) with PTFs.
- z/VM V5.4 with PTFs.
- z/VSE V4.2 with PTFs.
- z/TPF V1.1 PUT 4 with APARs.
- Linux on System z distributions - for four ports per feature (#3367):
 - Novell SUSE SLES 10 SP2 and SLES 11.
 - Red Hat RHEL 5.2 and RHEL 6.
- Linux on System z distributions - for two ports per feature (#3369):
 - Novell SUSE SLES 10 and SLES 11.

- Red Hat RHEL 5 and RHEL 6.

CHPID type OSD without maximum port exploitation (one port per PCIe adapter is available for use; two ports per feature):

- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.

CHPID type OSE supporting 4 or 2 ports per feature:

- z/OS V1.8 (2).
- z/VM V5.4.
- z/VSE V4.2.

CHPID type OSM for intranode management network:

- z/OS V1.12 with PTFs.
- z/OS V1.11 with PTFs.
- z/OS V1.10 (1) with PTFs.
- z/VM V6.1 with PTFs.
- z/VM V5.4 with PTFs to define, modify, and delete OSM CHPID types when z/VM is the controlling LPAR for dynamic I/O.
- IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

CHPID type OSN supporting OSA-Express for NCP: Note: CHPID type OSN does not use ports. All communication is LPAR-to-LPAR.

- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 10 and SLES 11.
 - Red Hat RHEL 5 and RHEL 6.

Display OSAINFO for OSA-Express3 CHPID types OSD, OSM (exclusive to zEnterprise), and OSX on z114 and z10 requires at a minimum:

- z/OS V1.12 with PTFs.

Inbound workload queuing (IWQ) for OSA-Express3 CHPID types OSD and OSX on zEnterprise requires at a minimum:

- z/OS V1.12 with PTFs.
- z/VM V5.4 with PTFs for guest exploitation.

12x InfiniBand and 1x InfiniBand coupling links require at a minimum:

- z/OS V1.10 (1) with PTFs.
- z/OS V1.9 (2) with PTFs.
- z/OS V1.8 (3) with PTFs.
- z/VM V5.4 to define, modify, and delete an InfiniBand coupling link, CHPID type CIB, when z/VM is the controlling LPAR for dynamic I/O.

Coupling Facility Control Code Level 17

Exploiting the new functions of Coupling Facility Control Code Level 17, including support for up to 2047 Coupling Facility structures, requires at a minimum:

- z/OS V1.12 with PTFs.
- z/OS V1.11 with PTFs.
- z/OS V1.10 (1) with PTFs.
- z/VM V5.4 with PTFs for guest virtual coupling.

(1) z/OS V1.10 supports the zEnterprise System; however z/OS V1.10 support will be withdrawn September 30, 2011. After that date, the IBM z/OS Lifecycle Extension for z/OS V1.10 (5656-A01) is required for the zEnterprise System. Talk to your IBM representative for details. Certain functions and features of the zEnterprise System require later releases of z/OS . For the complete list of software support, see the PSP buckets and the [Software requirements](#) section. For more information on the IBM Lifecycle Extension for z/OS V1.10, see Software Announcement [211-002](#), dated February 15, 2011 .

(2) z/OS V1.9 support was withdrawn September 30, 2010. However, with the z/OS Lifecycle Extension (5646-A01), z/OS V1.9 supports the zEnterprise System. Talk to your IBM representative for details. No exploitation of new zEnterprise System functions is available with z/OS V1.9. For the complete list of software support, see the PSP buckets and the [Software requirements](#) section. For more information on the IBM Lifecycle Extension for z/OS V1.9 see Software Announcement [210-027](#), dated May 11, 2010 .

(3) z/OS V1.8 support was withdrawn September 30, 2009. However, with the z/OS Lifecycle Extension (5638-A01), z/OS V1.8 supports the zEnterprise 114. Talk to your IBM representative for details. No exploitation of new zEnterprise 114 functions is available with z/OS V1.8. Certain functions and features of the zEnterprise require later releases of z/OS . For the complete list of software support, see the PSP buckets and the [Software requirements](#) section. For more information on the IBM Lifecycle Extension for z/OS V1.8, refer to Withdrawal Announcement [910-149](#), dated June 15, 2010 .

Planning information

Customer responsibilities

Information on customer responsibilities for site preparation can be found in the Library section of Resource Link at

<http://www.ibm.com/servers/resourcelink>

Installability

The average installation time for a z114 is approximately:

- M05: 11 installer hours
- M10: 11.7 installer hours

This does not include planning hours. This assumes the Pre-Installation Configuration Service, a full System Assurance Product Review, and implementation of the cable services have been performed. See your IBM representative for details on these services.

Security, auditability, and control

The z114 uses the security and auditability features and functions of host hardware, host software, and application software.

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

IBM Electronic Services

IBM has transformed its delivery of hardware and software support services to help you achieve higher system availability. Electronic Services is a web-enabled solution that offers an exclusive, no-additional-charge enhancement to the service and support available for IBM servers. These services are designed to provide the opportunity for greater system availability with faster problem resolution and preemptive monitoring. Electronic Services comprises two separate, but complementary, elements: Electronic Services news page and Electronic Services Agent.

The Electronic Services news page is a single Internet entry point that replaces the multiple entry points traditionally used to access IBM Internet services and support. The news page enables you to gain easier access to IBM resources for assistance in resolving technical problems.

The Electronic Service Agent™ is no-additional-charge software that resides on your server. It monitors events and transmits system inventory information to IBM on a periodic, client-defined timetable. The Electronic Service Agent automatically reports hardware problems to IBM. Early knowledge about potential problems enables IBM to deliver proactive service that may result in higher system availability and performance. In addition, information collected through the Service Agent is made available to IBM service support representatives when they help answer your questions or diagnose problems. Installation and use of IBM Electronic Service Agent for problem reporting enables IBM to provide better support and service for your IBM server.

To learn how Electronic Services can work for you, visit

<http://www.ibm.com/support/electronic>

Terms and conditions

IBM Global Financing

Yes

Warranty period

One year.

An IBM part or feature installed during the initial installation of an IBM machine is subject to a full warranty effective on the date of installation of the machine. An IBM part or feature which replaces a previously installed part or feature assumes the remainder of the warranty period for the replaced part or feature. An IBM part or feature added to a machine without replacing a previously installed part or feature is subject to a full warranty effective on its date of installation. Unless specified otherwise, the warranty period, type of warranty service, and service level of a part or feature are the same as those for the machine in which it is installed.

Customer setup

No.

IBM Licensed Internal Code (LIC) is licensed for use by a customer on a Specific Machine, designated by serial number, under the terms and conditions of the IBM License Agreement for Machine Code, to enable a Specific Machine to function in accordance with its Specifications, and only for the capacity authorized by IBM and acquired by the customer.

You can obtain the agreement at

http://www.ibm.com/systems/support/machine_warranties/machine_code.html

or by contacting your IBM representative.

Warranty service

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

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

The following service is available as warranty for your machine type.

- 24 hours per day, 7 days a week, same day response

Maintenance services

The z114 continues the heritage in mainframe qualities of service with the proprietary maintenance package available for all customers on Warranty and Maintenance, per the agreement of the maintenance contract. Once the maintenance contract is terminated the Proprietary Maintenance Service will no longer be available and the level of IBM service offered will revert to the Base Service Offering. The Base Service Offering will still provide the required direction for a service provider to repair a machine but will not include IBM intellectual property of step-by-step instructions, videos, and graphics.

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

The following service selections are available as maintenance options for your machine type.

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

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

The following service is provided.

- 24 hours per day, 7 days a week, same day response

Warranty service upgrades

Usage plan machine

No

IBM hourly service rate classification

Three

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

Field-installable features

Yes

Model conversions

Yes

Machine installation

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

Graduated program license charges apply

No

Licensed internal code

IBM Licensed Internal Code (LIC) is licensed for use by a customer on a specific machine, designated by serial number, under the terms and conditions of the IBM License Agreement for Machine Code, to enable a specific machine to function in accordance with its specifications, and only for the capacity authorized by IBM and acquired by the customer. You can obtain the agreement at

http://www.ibm.com/systems/support/machine_warranties/machine_code.html

or by contacting your IBM representative.

For machine type-models 2818-M15, 2818-M32, 2818-M49, 2818-M66, and 2818-M80, Elliptical Curve Cryptography (ECC) technology is included with the IBM zEnterprise 114 cryptography features. This technology is delivered through the machine's Licensed Internal Code, and requires license terms in addition to the standard IBM License Agreement for Machine Code (LMC) referenced above. These additional terms are delivered through the LMC's Addendum for Elliptical Curve Cryptography. This ECC Addendum is delivered with the machine along with the LMC when a cryptography feature is included in the IBM zEnterprise 114 order, or when a cryptography feature is carried forward as part of an MES order into IBM zEnterprise 114 and is available for review by contacting an IBM representative.

Pricing

For all local charges, contact your IBM representative.

Model conversions

Model
From To

R07 M05
R07 M10
S07 M05
S07 M10

E10 M05
E10 M10

M05 M10
M10 M15

Parts removed or replaced become the property of IBM and must be returned.

Feature conversions

Contact your IBM representative for feature conversion prices.

IBM Global Financing

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

<http://www.ibm.com/financing>

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

Order now

To order, contact the Americas Call Centers or your local IBM representative, or your IBM Business Partner.

To identify your local IBM representative or IBM Business Partner, call 800-IBM-4YOU (426-4968).

Phone: 800-IBM-CALL (426-2255)
Fax: 800-2IBM-FAX (242-6329)
Internet: callserv@ca.ibm.com
Mail: IBM Teleweb Customer Support
ibm.com® Sales Execution Center, Americas North
3500 Steeles Ave. East, Tower 3/4
Markham, Ontario
Canada
L3R 2Z1

Reference: YE001

The Americas Call Centers, our national direct marketing organization, can add your name to the mailing list for catalogs of IBM products.

Note: Shipments will begin after the planned availability date.

Trademarks

zEnterprise, z/VSE, z10, PR/SM, PowerVM, Hypervisor, HiperSockets, z9, Processor Resource/Systems Manager, MVS, DFSMSHsm, IMS, VSE/ESA and Electronic Service

Agent are trademarks of IBM Corporation in the United States, other countries, or both.

IBM, z/OS, z/VM, System z, AIX, System x, BladeCenter, System z10, DB2, WebSphere, DataPower, POWER7, ESCON, System p, FICON, System z9, z/Architecture, zSeries, Express, Resource Link, System Storage, Parallel Sysplex, Sysplex Timer, RETAIN, CICS, VTAM, Tivoli, Redbooks and ibm.com are registered trademarks of IBM Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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

Terms of use

IBM products and services which are announced and available in your country can be ordered under the applicable standard agreements, terms, conditions, and prices in effect at the time. IBM reserves the right to modify or withdraw this announcement at any time without notice. This announcement is provided for your information only. Additional terms of use are located at

<http://www.ibm.com/legal/us/en/>

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

<http://www.ibm.com/planetwide/us/>

Corrections

(Corrected on June 24, 2013)

Feature names were revised.

(Corrected on August 4, 2011)

An item was added to the "Planned availability date" section, and revisions were made under "HMC system support" and "Pricing."