Greg Daynes – IBM STSM - z/OS Installation and Deployment Architect
5 August 2010

z/OS Software Positioning for the IBM zEnterprise System

Permission is granted to SHARE Inc. to publish this presentation paper in the SHARE Inc. proceedings; IBM retains the right to distribute copies of this presentation to whomever it chooses.

© 2010 IBM Corporation

Agenda

- IBM zEnterprise 196 Overview
- z196 Support by z/OS Releases
- Migration Considerations
  - –General
  - –Migration Considerations for Selected Functions
  - –Multisystem/Sysplex Considerations
  - –Exploitation Considerations for Selected Functions
- Migration Paths
  - –General Recommendations and Considerations
  - –Recommended Paths for Supported z/OS Releases
- Summary
- Backup
  - –Cryptographic Support

© 2010 IBM Corporation

z/OS Software Positioning for the IBM zEnterprise System

The IBM zEnterprise System is comprised of the IBM zEnterprise 196, the IBM zEnterprise Unified Resource Manager, built-in support for the IBM zEnterprise BladeCenter Extension (zBX), and optimizers or IBM blades.

Preparing z/OS systems to support the new IBM zEnterprise System is easy. All supported releases can run on the new server; and enhancements in SMP/E V3.5 integrate the identification, verification, acquisition, and installation of required PTFs into standard SMP/E processing.

In this session, the speaker will cover z/OS software and configuration changes required to support the latest IBM zEnterprise System, the IBM zEnterprise 196 central processor complex (CPC), the IBM zEnterprise Unified Resource Manager, and the IBM zEnterprise BladeCenter Extension (zBX). He will describe the different functionality supported on each z/OS release, as well as any migration or exploitation actions required to use new functions on the IBM zEnterprise System. Systems programmers and their managers who will upgrade to the latest server will find lots of useful information in this session.

z/OS Software Positioning for the IBM zEnterprise System

Scope

- **The IBM zEnterprise System is comprised of:**
  - –the IBM zEnterprise 196 (z196),
  - –the IBM zEnterprise Unified Resource Manager,
  - –the IBM zEnterprise BladeCenter Extension (zBX) with
    - Optimizers or IBM blades.
- **z/OS platform software requirements for:**
  - –Hardware upgrade to z196
    - With or without zBX
  - –Install a new z196 server
    - With or without zBX
- **Outside scope**
  - –z/VM (native), zLinux, and zTPF considerations
  - –Non-z/OS software required for zBX solutions

3 © 2010 IBM Corporation

z/OS Software Positioning for the IBM zEnterprise System

IBM zEnterprise System Functions and Features (GA)

Five hardware models
Quad-core 5.2 GHz processor chips
Up to 80 processors configurable as CPs, zAAPs, zIIPs, IFLs, ICFs, or optional SAPs
Increased capacity processors
Out of order instruction execution
Improved processor cache design
Up to 15 subcapacity CPs at capacity settings 4, 5, or 6
Up to 3 TB real memory (1 TB per LPAR)
Improved availability with Redundant Array of Independent Memory (RAIM)
Power save functions
8 slot, 2 domain I/O drawer
On Demand enhancements
Crypto Express3 enhancements
IBM zEnterprise Unified Resource Manager (from HMC)
New and enhanced instructions
6.0 GB/sec InfiniBand I/O interut

(z/OS Support in blue)

Capacity Provisioning enhancements
2 New OSA CHPIDs – OSX and OSM
Three subchannel sets per LCSS
FICON Discovery and AutoConfiguration (zDAC)
Concurrent I/O drawer add, remove, replace
Doubled number of HiperSockets to 32
CFCC Level 17 enhancements
Up to 80 External Coupling Link Ports
Up to 128 Coupling Link CHPIDs Defined
Optional water cooling
Optional High Voltage DC power
Optional overhead I/O cable exit
STP enhancements
zBX-002 Smart Analytics Optimizer
zBX-002 IBM Blades

4 © 2010 IBM Corporation

IBM zEnterprise System Overview

IBM zEnterprise System is a first-of-a-kind workload-optimized multiplatform technology offering. The system consists of the IBM zEnterprise 196 (z196) central processor complex (CPC), the IBM zEnterprise Unified Resource Manager, and built-in support for the IBM zEnterprise BladeCenter Extension (zBX) Model 002. The IBM zEnterprise 196 is designed with improved scalability, performance, security, resiliency, availability, and virtualization. The z196 Model M80 provides up to 1.6 times the total system capacity of the z10 EC Model E64, and all z196 models provide up to twice the available memory of the z10 EC. The zBX

infrastructure works with the IBM zEnterprise 196 to enhance System z virtualization and management to deliver an integrated hardware platform that spans System z mainframe and POWER7 technologies. The IBM zEnterprise Unified Resource Manager, delivered with the z196, is designed to deliver end-to-end virtualization and management along with the ability to optimize technology deployment according to individual workload requirements.

zEnterprise System is designed to deliver:

- **Improved total system capacity** in a 96-way core design.
  - Improved performance of all workloads
  - Massive scalability for secure data serving and transaction processing and large-scale consolidation
  - Up to 80 cores that are customer-configurable
- **Quad-core 5.2 GHz processor chips** with 100+ new instructions that enable improved code efficiency, and are designed to help improve the execution of CPU-intensive workloads.
- **Up to 3 terabytes (TB) of available real memory per server** for growing application needs (with up to 1 TB real memory per LPAR).
- Improved availability in the memory subsystem with **redundant array of independent memory (RAIM)**.
- Availability improvement - **concurrent add/delete of I/O drawer**.
- **Increased scalability** with 45 available subcapacity settings.
- **Crypto Express3 cryptographic enhancements**.
- **Auto-renewal of On/Off Capacity on Demand records** which is designed to eliminate the need to manually replenish your On/Off CoD records.
- **Purchase of unassigned CP or IFL capacity** via Customer Initiated Upgrade.
- **Up to 128 coupling CHPIDs** (double that of z10) and up to 80 physical external coupling links (12x InfiniBand, 1x InfiniBand, ISC-3), a 25% increase compared to z10.
- Up to 72 I/O, networking, and crypto features.
- Up to 240 ESCON channels.
- **Up to 32 HiperSockets** (double the number available on z10 EC).
- **Energy efficiencies** - more performance to help you run more efficiently in the same footprint.
  - Energy monitoring capability and power save option for the processor, which is designed to provide a 25 - 30% reduction in power consumption depending on the system configuration
  - Optional water cooling, providing the ability to cool systems with user-chilled water
  - Optional high-voltage DC power, which can save System z customers 1-3%, on average, on their power bills
- **Top exit I/O cabling**, which is designed to provide increased flexibility and increase air flow in raised-floor environments.
- **Unified Resource Manager**, enabling management and virtualization of heterogeneous workloads. Unified Resource Manager manages the deployment of heterogeneous hardware resources based on individual workload requirements by providing:
  - Performance management
  - Integrated private data network
  - Virtual server lifecycle management
  - Hypervisor management
  - Integrated operational controls
  - Management as system z firmware

Planned availability dates

September 10, 2010

- Features and functions for the z196
- z196 Models M15, M32, M49, M66, and M80
- Water Cooling (#0159) for z196
- Manage suite (#0019) for z196
- z9 EC upgrades to air-cooled z196
- z9 EC upgrades to water-cooled z196

© 2010 IBM Corporation

- z10 EC upgrades to air-cooled z196
- z10 EC upgrades to water-cooled z196
- 3-in-1 Bolt Down Kit (#8008 or #8009) for new build z196

November 9, 2010
- HMC with Dual Ethernet (#0091) on z10 EC and z10 BC
- TKE workstation (#0841) on z10 EC and z10 BC
- TKE 7.0 LIC (#0860) on z10 EC and z10 BC

November 19, 2010
- Manage suite enhancement (#0019) functions for z196
- Automate suite (#0020) for z196
- IBM Smart Analytics Optimizer for DB2 for z/OS for z196 in blade quantities of 7, 14, or 28 (#0019 or #0020)
- POWER7 blade for z196 in a zBX (#0019 or #0020)

3Q2010
- Inbound workload queuing (IWQ) for OSA-Express3 on z10

December 17, 2010
- IBM Smart Analytics Optimizer enablement (#0251) on z10 EC and z10 BC
- IBM Smart Analytics Optimizer for DB2 for z/OS for z196 in blade quantities of 42 or 56 (#0019 or #0020)

December 31, 2010
- 3-in-1 Bolt Down Kit (#8008 or #8009) MES z196
- MES features for Models M15, M32, M49, M66, and M80
- Model conversions for z196

March 17, 2011
- z10 EC with zBX Model 001 upgrades to z196 with zBX Model 002

© 2010 IBM Corporation

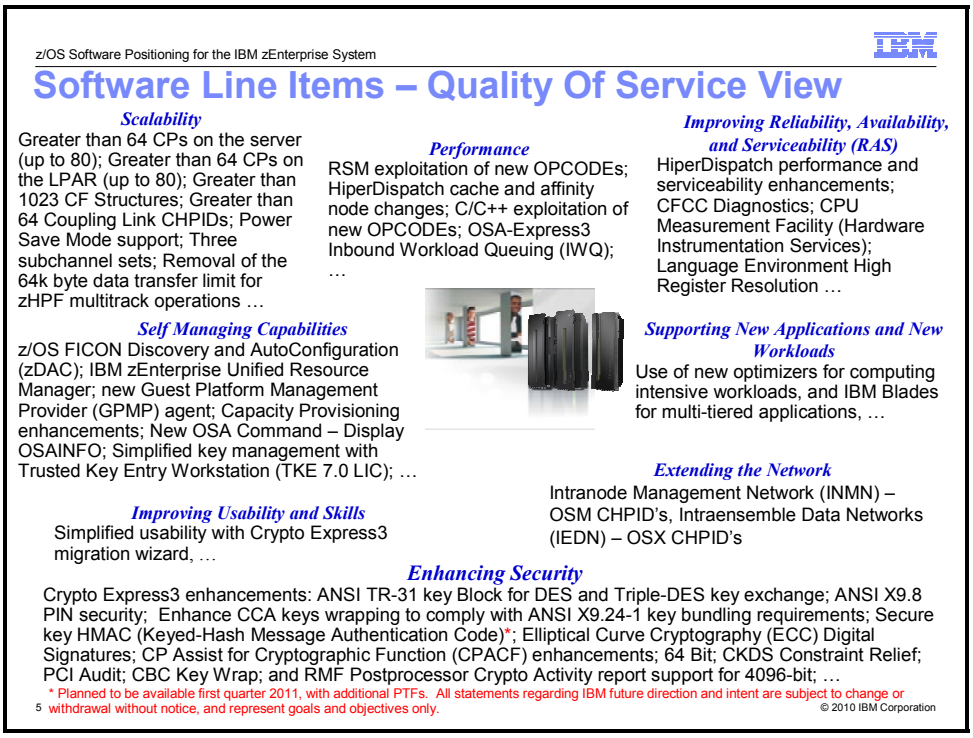

z/OS Software Positioning for the IBM zEnterprise System

Scalability and Performance Enhancements

The IBM zEnterprise 196 server adds additional scalability and performance capabilities for your z/OS environment.

- The new 96-way core design (with 80 cores that are customer configurable) delivers massive scalability for secure data serving and transaction processing for large-scale businesses. The performance of a z196 (2817) processor is expected to be 1.3 to 1.5 times the performance of a z10 EC (2097) based on workload and model. The largest z196 (z196 Model M80) is expected to exceed 1.6 times the capacity of the largest z10 (2097-764). It has up to twice the available real memory, 3 terabytes (TB) per server (with up to 1 TB real memory per LPAR) compared to the z10 EC Model E64. New quad-core 5.2 GHz processor chips, with more than 100 new instructions to enable improved code efficiency, are also designed to help improve the execution of Java and CPU-intensive workloads. For example, it is anticipated that z/OS and zEnterprise 196 servers can provide a significant performance improvement for Java workloads.
- In addition, you now have the ability to extend the amount of addressable storage capacity to help facilitate storage growth with the introduction of a third subchannel set - an additional 64K subchannels - to help complement other functions such as "large" or extended address volumes (EAV) and HyperPAV. This may also help facilitate consistent device address definitions, simplifying addressing schemes for congruous devices. The first subchannel set (SS0) allows definitions of any type of device (such as bases, aliases, secondaries, and devices that do not implement the concept of associated aliases or secondaries). The second and third subchannel sets (SS1 and SS2) can now both be used for disk alias devices (both primary and secondary devices) and/or Metro Mirror secondary devices only. The third subchannel set supports ESCON, FICON and zHPF protocols.
- The zEnterprise 196 provides many enhancements to a z/OS Parallel Sysplex environment:
  - Connectivity improvements with up to 80 coupling links: The IBM zEnterprise 196 increases the number of external coupling links allowed from 64 to 80. This allows the full configuration of 32 PSIFB links and 48 ISC-3 links to be used. In addition, you 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 coupling 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 is also required. With z196, the number of coupling CHPIDs per server has been increased 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.
    - 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. z196 and CFCC level 17 also provide improved serviceability of Coupling Facilities with enhanced data collection and triggering of nondisruptive CF dumps.
    - z/OS on z196 servers with CFCC Level 17 also supports more connectors to list and lock structures. XES and CFCC already support 255 connectors to cache structures. With this new support XES also supports up to 247 connectors to a lock structure, 127 connectors to a serialized list structure, and 255 connectors to an unserialized list structure.
- z/OS V1.12 on z196 servers is designed to provide improved networking performance with OSA-Express3 in QDIO mode with inbound workload queuing (IWQ). IWQ creates multiple input queues and allows OSA to differentiate workloads "off the wire" and then assign work to specific input queues 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 use appropriate processing resources for each input queue. This approach allows multiple concurrent z/OS processing threads to process each unique input queue, 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 is expected to result in improved overall system performance and scalability.
    - It is anticipated that networking performance for interactive workloads can be improved significantly, depending on amount of data being transferred, presence of bulk-data traffic in the mix, and whether communication is z/OS to z/OS, or z/OS to distributed system. For example, interactive networking response time improvements of 30-50% on a System z10 model 2097-E64 were recorded.
    - It is also anticipated that bulk-data (streaming) workloads can also benefit with OSA-Express3 IWQ and its ability to reduce the amount of costly network retransmissions (by reducing the incidence of out-of-order packets). A streamlined Communications Server execution path for Sysplex Distributor over IWQ is expected to improve performance for sysplex-distributed traffic as well.
    - When inbound workload queuing is enabled for a QDIO interface, inbound streaming bulk data is processed on an ancillary input queue (AIQ). This function is expected to improve throughput while reducing CPU consumption for inbound streaming bulk data.
    - When inbound workload queuing is enabled for a QDIO interface, inbound Sysplex Distributor traffic is processed on an AIQ. This function is expected to improve performance for inbound Sysplex Distributor traffic that is routed to a target stack.
- z/OS V1.12 XL C/C++ also exploits new instructions in the IBM zEnterprise System 196 server. The z/OS V1.12 XL C/C++ compiler provides new ARCHITECTURE(9) and TUNE(9) options to help you exploit new instructions that are available on z196 servers. This is designed to provide better performing applications tuned for the new server. Additional optimization and tuning have been made to improve the floating-point performance. These changes can improve the performance of generated code without the need for changes to the source code. A performance improvement of over 11% was observed using code generated by the z/OS V1.12 XL C/C++ compiler with high optimization when compared to code generated using the z/OS V1.11 XL C/C++ compiler. Performance improvements are based on internal IBM lab measurements using the ILP32, XPLINK, ARCH(9), TUNE(9), HGPR, O3, HOT, and IPA(LEVEL(2)) with PDF compiler options. Performance results for specific applications will vary; some factors affecting performance are the source code and the compiler options specified.

Expanding the Networking Capabilities

- The new zEnterprise System unifies the zEnterprise 196 server and the zEnterprise BladeCenter Extension, through the use of two new "internal" networks. Configured through OSA-Express3 adapters, the intranode management network (INMN) is used to provide system management functions and the intraensemble data network (IEDN) can be used to send data between the systems.

  - The intranode management network (INMN) is a private and physically isolated OSA-Express3 1000BASE-T Ethernet internal management network operating at 1 Gbps, that connects all zEnterprise System resources (CPCs, Blade Centers, 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 your application data traffic. The INMN is exclusive to zEnterprise System, and is configured with a CHPID type of OSM.
  - Intraensemble data network (IEDN) is a private OSA-Express3 10 Gigabit Ethernet 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 customer-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 zEnterprise System, and is configured with a CHPID type of OSX.
- The z/OS Communications Server also eases the burden of configuration for these new OSA-Express3 CHPID types by dynamically finding and activating up to two OSA-Express3 adapters connected to the Intranode Management Network, and requiring minimal configuration for OSA-Express3 adapters connected to the Intraensemble Data Network.

System Management Enhancements

- A new IBM System z Discovery and AutoConfiguration (zDAC) function, available with z/OS V1.12 on zEnterprise 196 servers, can help simplify the configuration of new and changed disk and tape I/O devices. zDAC can save time by discovering new and changed fabric-attached devices, and suggesting configurations aligned with best practices for availability and with the I/O policies that you set. For example, with zDAC, building an updated I/O configuration can take as little as a few minutes. zDAC is designed to perform discovery for a single system or for all the systems in a sysplex. These capabilities are integrated with HCD and z/OS Hardware Configuration Manager (HCM). 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), HyperPAV, DCM specifications, and preferred control unit and device number ranges.
- The new zEnterprise System also provides end-to-end workload monitoring and other systems management capabilities for System z ensembles, through the new IBM zEnterprise Unified Resource Manager. System z ensembles are collections of one or more zEnterprise System nodes in which each node comprises a z196 server and its optionally attached IBM zEnterprise BladeCenter Extension (zBX) model 002. An ensemble can consist of a single z196 server running z/OS images and z/VM hosting Linux for System z images but without a zBX attached, or it can consist of from 1 to 8 z196 servers, at least one of which has a zBX attached. The resources of a zEnterprise System ensemble are managed and virtualized as a single pool, integrating system and workload management across the multisystem, multitier, multiarchitecture environment. The zEnterprise Unified Resource Manager uses the intranode management network (INMN) for communication.
- z/OS integrates with this new management environment. A new agent, Guest Platform Management Provider (GPMP), communicates with z/OS WLM and provides basic data (such as system resource utilization, system delays, and paging delays) back to the zEnterprise Unified Resource Manager over the INMN network. The zEnterprise Unified Resource Manager can add additional workload relationships from the ensemble components to your z/OS workload; for example, linking a transaction that started on the zBX back to DB2 on z/OS data.
- In addition to IWQ, OSA-Express3 also 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 can help you improve the overall management, serviceability, and usability of OSA-Express3. The display OSAINFO requires OSA-Express3 CHPID types OSD, OSM, and OSX, and z/OS V1.12.

© 2010 IBM Corporation

Security Enhancements

- The z/OS security enhancements in support of the z196 server are provided in the Cryptographic Support for z/OS V1.10 through z/OS V1.12 Web deliverable which is planned to be made available September 10, 2010. To obtain this Web deliverable, when available, visit: **http://www.ibm.com/systems/z/os/zos/downloads/.** This Web deliverable is planned to support z/OS V1.10, z/OS V1.11, and z/OS V1.12. A number of ICSF enhancements are planned:
  - An improved symmetric key store (CKDS), to significantly increase the number of keys that can be stored in the CKDS and to reduce the overhead associated with updates to the CKDS. These improvements are intended to help position ICSF as the enterprise key store and reduce the batch update window for customers that add or rotate keys in the CKDS frequently.
  - 64-Bit enablement for the full set of ICSF callable services. This support is designed to reduce application overhead and simplify the use of ICSF callable services in 64-bit environments by eliminating the need to switch between 64-bit and 31-bit modes.
  - Additional content in audit records to include identity information intended to help you meet the requirements of specialized industry standards such as PCI DSS. New sections are planned to be added to SMF type 82 records to contain the identities of the ICSF server and cryptographic service caller. This enhancement is intended to help simplify compliance audits by minimizing the need to correlate ICSF audit records with other security audit records in order to obtain the identity information associated with the ICSF audit event.
  - API support for new symmetric encryption modes introduced in the IBM zEnterprise 196 (z196) server CP Assist for Cryptographic Functions (CPACF). The new encryption modes include CMAC, CCM, GCM, CFB, OFB, XTS and CBC-S. These new modes will be supported for both CPACF clear keys and CPACF protected keys. This new support is intended to help application programmers take advantage of these new encryption modes while using the familiar CCA callable services provided by z/OS ICSF.

  ICSF is also planned to provide new support for new functions on zEnterprise 196 servers with the Crypto Express3 feature and minimum driver and microcode levels dated August 2010 or later:
  - New CCA key token wrapping mechanisms designed to bring CCA key tokens into compliance with banking and finance industry standards that mandate key bundle compliance. This support is intended to enable ICSF to utilize key tokens that comply with ANSI x9.24-1 key management requirements for symmetric keys wrapped under either symmetric master keys or key-encrypting keys.
  - Support for enhanced Personal Identification Number (PIN) protection techniques for online PIN verification via ATM and Point of Sale (POS) systems. The enhancements are are intended to provide the functions defined as extensions to the ANSI x9.8 and ISO 9564-1 PIN Management and Security standards.
  - New CCA functions implementing the Elliptic Curve Digital Signature Algorithm (ECDSA). These are intended to allow applications using ICSF to generate elliptic curve key pairs, store elliptic curve keys in the PKDS, and both generate and verify elliptic curve digital signatures. Support is planned for both NIST and Brainpool standard curves, and to support the Crypto Express3 Coprocessor card performing ECDSA within the bounds of the secure hardware.
  - New CCA functions for Hashed Message Authentication Codes (HMAC) using the SHA-1 and SHA-2 families of hash functions. HMACs are widely used in internet and storage protocols and are used in specific geographies for the banking and finance industry. This support is designed to allow applications using ICSF to generate HMAC keys, store HMAC keys in the CKDS, and to generate and verify HMACs. This support is to be made available with the Cryptographic Support for z/OS V1.10 through z/OS V1.12 Web deliverable with the PTF for APAR OA33260, planned for February 2011 availability.

© 2010 IBM Corporation

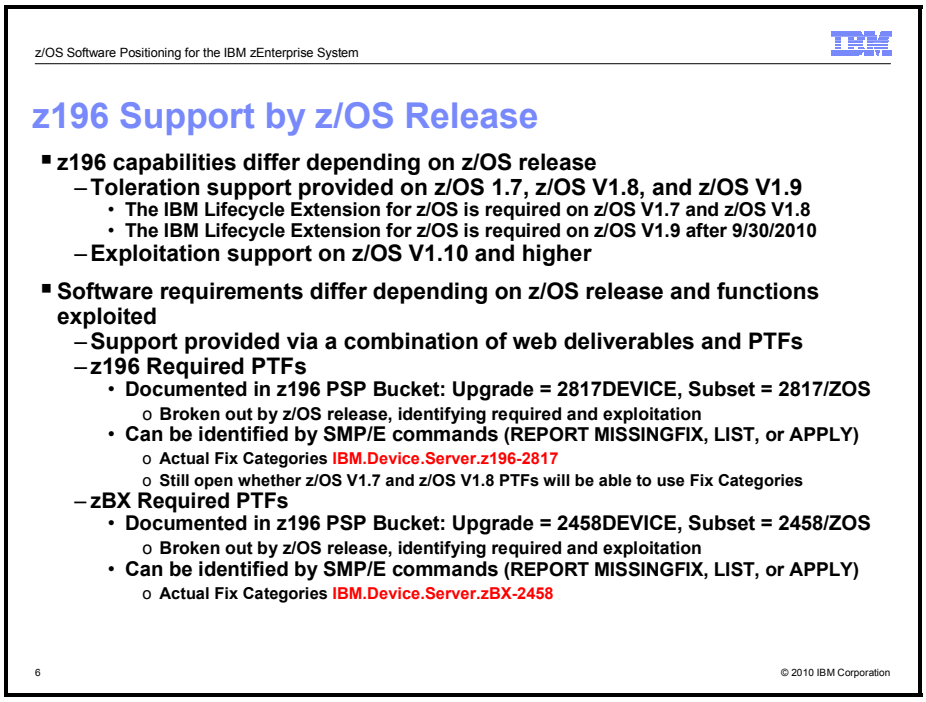

z196 Support by z/OS Release

The software support for the z196 servers, similar to the IBM System z10 EC and z10 BC servers, differs by z/OS release level, and is provided by a combination of FMIDs and PTFs. The specific PTFs that need to be installed will depend on the z/OS release. In addition, exploitation of certain functions, specifically some cryptographic functions, requires the installation of an FMID. These FMIDs are downloadable Web Deliverables, which are unpriced and can be downloaded at: http://www.ibm.com/eserver/zseries/zos/downloads.

For a z196 CPC, PTFs are identified in the 2817DEVICE PSP bucket (Subset 2817/ZOS). For an IBM zEnterprise BladeCenter Extension (zBX) attached to your z196 CPC, either to support IBM Smart Analytics Optimizer for DB2 for z/OS, V1.1 (5697-AQT) or to support IBM POWER7 Blades, the PTFs are identified in the 2458DEVICE PSP bucket (Subset 2458/ZOS). In each PSP bucket, the content is dependent on the z/OS release you will run on the z196 server. If you reviewed the PSP buckets some time ago, review them again to ensure that any newly identified z/OS service has been installed. To assist you in determining if you have the recommended service (identified in these PSP buckets) installed on your system, you can use the SMP/E REPORT MISSINGFIX command in conjunction with the FIXCAT type of HOLDDATA, as follows:

a. Acquire and RECEIVE the latest HOLDDATA onto your z/OS system(s). Use your normal service acquisition portals or download the two (2) year HOLDDATA directly from http://service.software.ibm.com/holdata/390holddata.html. Ensure you select **Full** from the Download NOW column (last 730 days) to receive the FIXCAT HOLDDATA, as the other files do not contain FIXCAT HOLDDATA.
b. Run the SMP/E REPORT MISSINGFIX command on your z/OS systems and specify one or more of the following Fix Categories (FIXCAT):
   - IBM.Device.Server.z196-2817
   - IBM.Device.Server.z196-2817.ParallelSysplexInfiniBandCoupling
   - IBM.Device.Server.z196-2817.ServerTimeProtocol
   - IBM.Device.Server.z196-2817.zHighPerformanceFICON
   - IBM.Device.Server.zBX-2458
   - IBM.Device.Server.zBX-2458.ISAOPT

   The report will identify any missing coexistence and fallback PTFs for that system. For complete information about the REPORT MISSINGFIX command, see *SMP/E Commands*.
c. Periodically, you might want to acquire the latest HOLDDATA and rerun the REPORT MISSINGFIX command to find out if there are any new PTFs recommended for the z196 server.

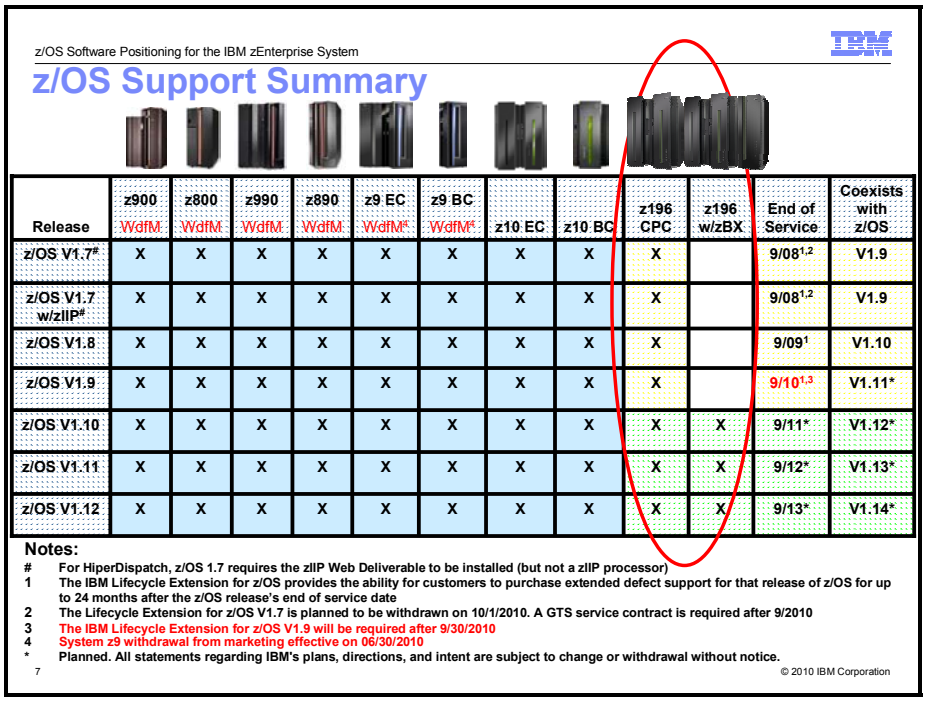

z/OS Software Positioning for the IBM zEnterprise System

z/OS Support Summary

Release	z900 WdfM	z800 WdfM	z990 WdfM	z890 WdfM	z9 EC WdfM[4]	z9 BC WdfM[4]	z10 EC	z10 BC	z196 CPC	z196 w/zBX	End of Service	Coexists with z/OS
z/OS V1.7#	X	X	X	X	X	X	X	X	X		9/08[1,2]	V1.9
z/OS V1.7 w/zIIP#	X	X	X	X	X	X	X	X	X		9/08[1,2]	V1.9
z/OS V1.8	X	X	X	X	X	X	X	X	X		9/09[1]	V1.10
z/OS V1.9	X	X	X	X	X	X	X	X	X		9/10[1,3]	V1.11*
z/OS V1.10	X	X	X	X	X	X	X	X	X	X	9/11*	V1.12*
z/OS V1.11	X	X	X	X	X	X	X	X	X	X	9/12*	V1.13*
z/OS V1.12	X	X	X	X	X	X	X	X	X	X	9/13*	V1.14*

Notes:

#	For HiperDispatch, z/OS 1.7 requires the zIIP Web Deliverable to be installed (but not a zIIP processor)
1	The IBM Lifecycle Extension for z/OS provides the ability for customers to purchase extended defect support for that release of z/OS for up to 24 months after the z/OS release's end of service date
2	The Lifecycle Extension for z/OS V1.7 is planned to be withdrawn on 10/1/2010. A GTS service contract is required after 9/2010
3	The IBM Lifecycle Extension for z/OS V1.9 will be required after 9/30/2010
4	System z9 withdrawal from marketing effective on 06/30/2010
*	Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

7 © 2010 IBM Corporation

z/OS Support Summary

All supported z/OS releases can run on a z196 server. Similarly, all supported z/OS releases can participate in a sysplex that has a CF or operating system image on a z196 server. Not all System z196 functions are available in every z/OS release. Other than the minimum support, there is no requirement to install either software or hardware first to exploit a specific function.

IBM zEnterprise 196 requires at a minimum:

- z/OS V1.10 or higher with PTFs.
- z/OS V1.9[1] with PTFs.
- z/OS V1.8[2] with the IBM Lifecycle Extension for z/OS V1.8 (5638-A01) with required maintenance.
- z/OS V1.7[3] with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with required maintenance.

IBM zEnterprise Unified Resource Manager and IBM zEnterprise BladeCenter Extension (zBX) model 002 requires at a minimum:

- z/OS V1.10 with PTFs.

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

[1] z/OS V1.9 supports zEnterprise System, however, z/OS V1.9 support will be withdrawn September 30, 2010. After that date, the IBM z/OS Lifecycle Extension for z/OS V1.9 (5646-A01) is required for zEnterprise System. Talk to your IBM representative for details. No exploitation of new zEnterprise System functions is available with z/OS V1.9. Certain functions and features of the zEnterprise System require later releases of z/OS. For more information on the IBM Lifecycle Extension for z/OS V1.9 see Software Announcement 210-027 dated May 11, 2010.

[2] 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 System. Talk to your IBM representative for details. No exploitation of new zEnterprise System functions is available with z/OS V1.8. Certain functions and features of the zEnterprise System require later releases of z/OS. For more information on the IBM Lifecycle Extension for z/OS V1.8 see Software Announcement 209-180 (RFA 53080) dated June 9, 2009.

[3] z/OS V1.7 support was withdrawn September 30, 2008. However, with the z/OS Lifecycle Extension (5637-A01), z/OS V1.7 supports the zEnterprise System. No exploitation of new zEnterprise System functions is available with z/OS V1.7. Certain functions and features of the zEnterprise System require later releases of z/OS. The z/OS Lifecycle Extension for z/OS V1.7 provided defect support until September 30, 2010. Contact your IBM representative if you need support passed that date. For more information on the IBM Lifecycle Extension for z/OS V1.7 see Software Announcement 208-283 (RFA 47640) dated August 12, 2008.

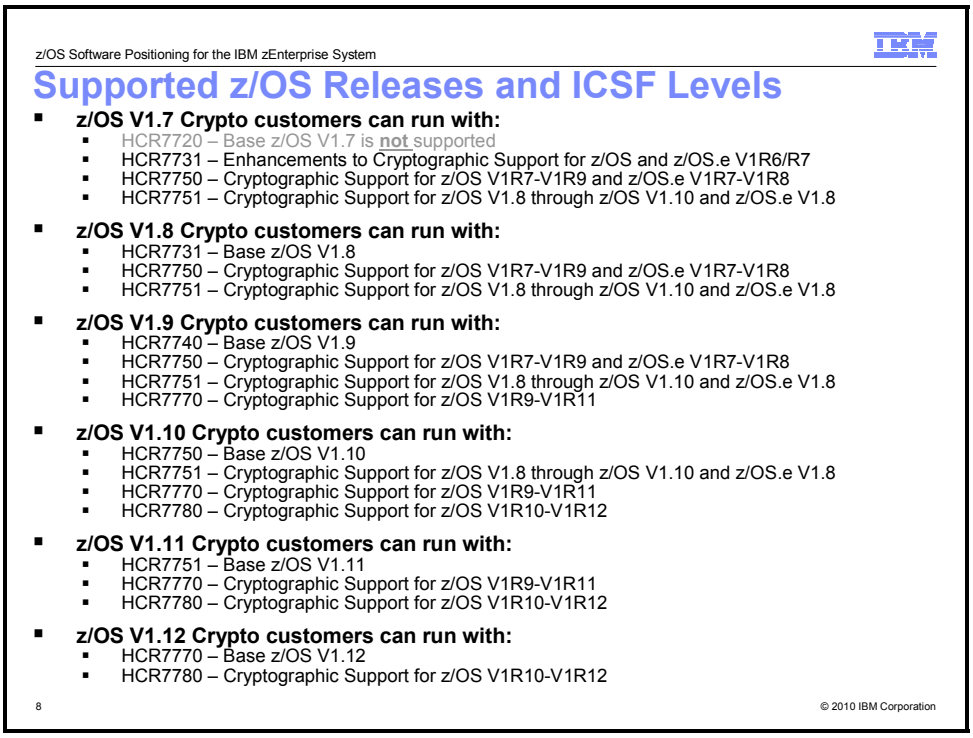

Supported z/OS Releases and ICSF Levels

The level of function provided for cryptographic toleration differs by z/OS release and the ICSF web deliverable that is installed. For z/OS V1.10 and higher, exploitation of z196 cryptographic support is provided by Cryptographic Support for z/OS V1R10-V1R12 (FMID HCR7780). Note that this level of ICSF is NOT integrated in z/OS V1.12 and will need to be downloaded and installed even after ordering a z/OS V1.12 ServerPac.

a. Toleration PTFs are available for all z/OS releases, and all existing supported ICSF web deliverables.
b. (For z/OS V1.7) Minimally you will have to install toleration PTFs on top of the Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/R7 web deliverable (FMID HCR7731). The level of ICSF initially shipped with z/OS V1.7 (FMID HCR7720) does not run on a z196 server.
c. (For z/OS V1.7-z/OS V1.9) If you require Crypto Advanced Encryption Standard (AES) for 256-bit keys SHA-384 and 512 bit for message digest support, ISO Format 3 PIN blocks, support for RSA keys up to 4096 bits, Random Number Generator Long, or enhanced TKE Auditing, then you must install the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web Deliverable (FMID HCR7750).
d. (For z/OS V1.7-z/OS V1.10) If you require Support for 13- thru 19-digit Personal Account Numbers, ICSF Query Algorithms, or Key Token Policy support then you must install the Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8 Web deliverable (FMID HCR7751).
e. (For z/OS V1.9-z/OS V1.11) If you require protected key CP Assist for Cryptographic Function, new Crypto Express3 or Crypto Express3 -1P, then you must install the Cryptographic Support for z/OS V1R9-V1R11 Web deliverable (FMID HCR7770).
f. (For z/OS V1.10-z/OS V1.12) If you require X9.8 Pin, 64 Bit, z196 CPACF, HMAC, CKDS Constraint Relief, PCI Audit, ECC HW Support, CBC Key Wrap, then you must install the Cryptographic Support for z/OS V1R10-V1R12 Web deliverable (FMID HCR7780). The planned availability of this web deliverable is September 2010. The planned availability of HMAC support is first quarter 2011, with additional PTFs.

Note: While there are a lot of supported combinations of z/OS releases and ICSF supported levels, most likely you only have to be concerned with one (possibly two) bullets on the slide above. For more information regarding ICSF functionality available on each z/OS release, please see the “Backup Material” starting on page 62.

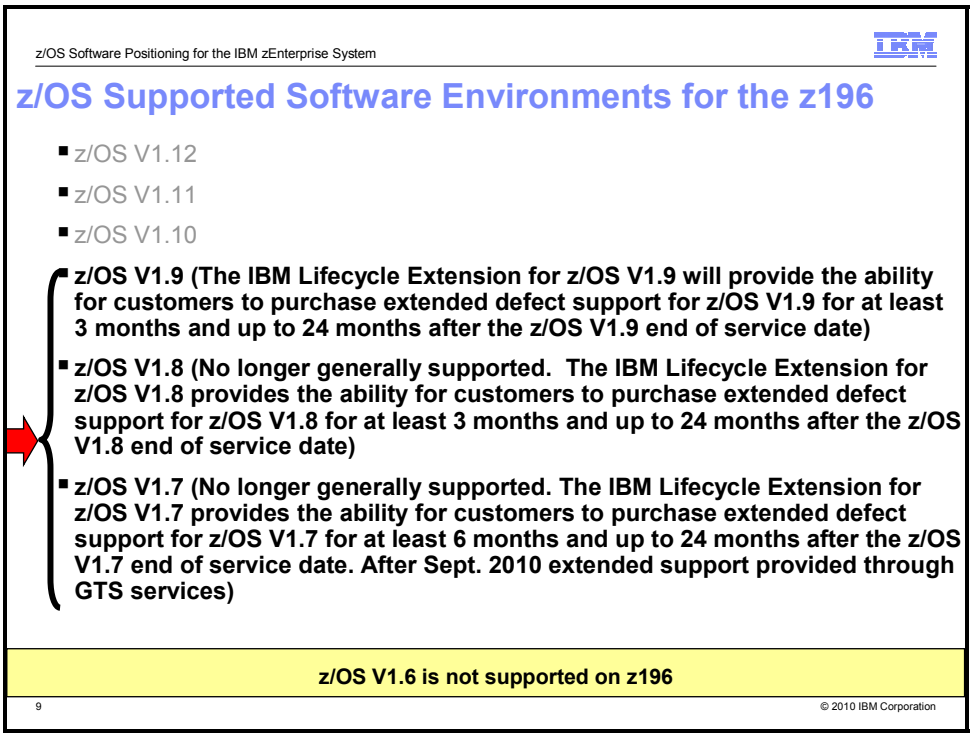

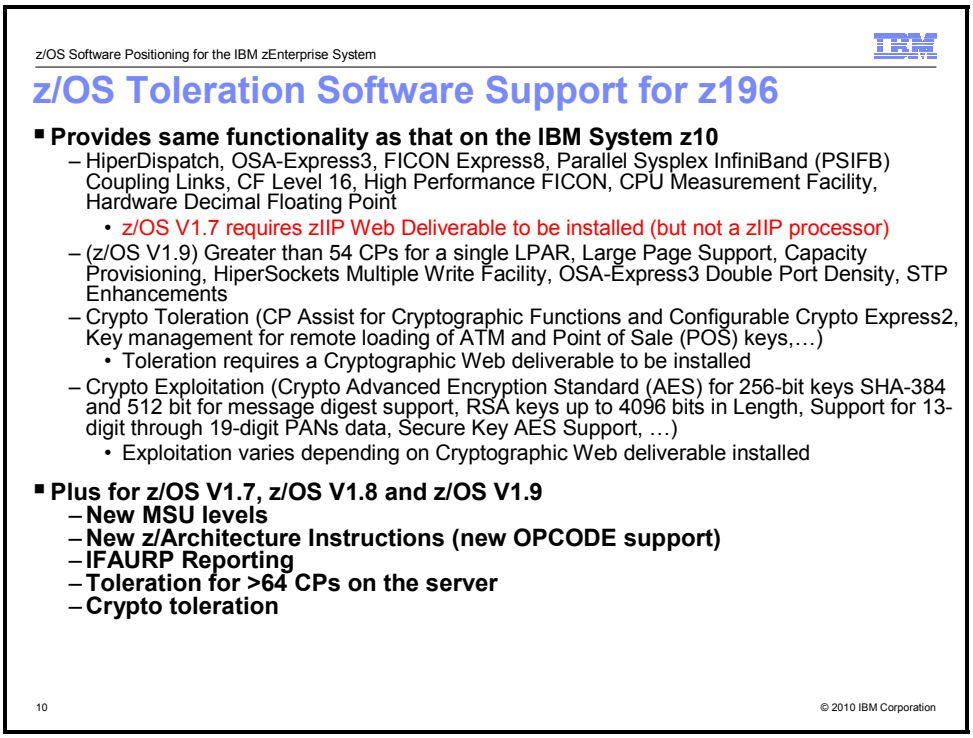

z/OS Toleration Support for a z196 Server

z/OS V1.7 – z/OS V1.9 Support is included for the following functions on the z196:

- z196 new instruction mnemonic support,

- HiperDispatch (z/OS V1.7 requires the IBM zIIP Support for z/OS V1R6-V1R7 and z/OS.e V1R6-V1R7 web deliverable to be installed, but not a zIIP processor),
- OSA-Express3 (GbE LX and SX, 1000BASE-T, 10 GbE LR and SR),
- FICON Express8,
- Parallel Sysplex InfiniBand (PSIFB) Coupling Links,
- High Performance FICON for System z (zHPF),
- Usage Report Program (IFAURP) support,
- toleration for >64 CPs on the server,
- (z/OS V1.8 and z/OS V1.9 only) OSA-Express3 CHPID OSD maximum ports,
- (z/OS V1.8 and z/OS V1.9 only) RMF Postprocessor Crypto Activity report - 4096-bit,
- (z/OS V1.8 and z/OS V1.9 only) enhancements to CPU Measurement Facility (HIS),
- Cryptographic toleration support.

The cryptographic toleration support is provided by PTFs installed on either level of ICSF originally shipped with the z/OS release (see prior chart) or PTFs installed on one of the web deliverables. The support varies based on the level of ICSF installed. See page 63 for the level of support provided in each web deliverable.

- Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/R7 web deliverable (HCR7731)
- Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 web deliverable (HCR7750)
- Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8 web deliverable (HCR7751)
- (z/OS V1.9 only) Cryptographic Support for z/OS V1R9-V1R11 web deliverable (HCR7770)

z/OS Toleration Software Support for z196

z/OS Software Positioning for the IBM zEnterprise System

Release	Base GA1 Support	OSA-Express3 (GbE LX and SX, 1000BASE-T, 10 GbE LR and SR)	New z/Architecture Instructions[3]	IB Coupling Links	Up to 128 Coupling Link CHPIDs	FICON Express8 (CHPID FC)	IFAURP Reporting	>64 CPs per Server	Crypto Toleration[4]	HiperDispatch	OSA-Express3 CHPID OSD maximum ports	RMF Postprocessor Crypto Activity report - 4096-bit[6]	CPU Measurement Facility (HIS)	> 64 CPs per LPAR	Cache and Affinity Node Changes	HiperDispatch Serviceability	LE High Register Resolution	Power Save Mode	CF Level 17	Extensions to zHPF Multitrack Operations	Three Subchannel Sets	Crypto Exploitation	Up to 32 HiperSockets	IBM zEnterprise BladeCenter Extension (zBX) support	Network Management	Performance Management	zDAC Support	New OSA Display Command	OSA-Express3 Inbound Workload Queuing (IWQ)	XL C/C++ ARCH(9) Tune(9)
	z196 PSP Bucket – 2817DEVICE 2817/ZOS																													
	Base Support																	Explicit z/OS Support												
z/OS V1.7[1]	P	P	P	P	P	P	P	P	WP	N	N	P[2]	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS V1.7[1] w/ZIIP	P	P	P	P	P	P	P	P	WP	WP	N	P[2]	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS V1.8[1]	P	P	P	P	P	P	P	P	P	BP	P	P	P	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS V1.9[2]	P	P	P	P	P	P	P	P	P	BP	P	P	P	N	N	N	N	P[5]	N	N	N	N	N	N	N	N	N	N	N	N

[1] – The Lifecycle Extension for z/OS V1.7 or z/OS V1.8 is required for support
[2] – The IBM Lifecycle Extension for z/OS V1.9 will be required after 9/30/2010
[3] – Does not include XL C/C++ support for ARCH(9) and TUNE(9) Options
[4] – A Crypto Web Deliverable is NOT required, but toleration PTF is needed even if a web deliverable is installed
[5] – Toleration PTFs needed to detect Power Saving Mode
[6] – Toleration PTFs recognize crypto control block, but doesn't provide updated reports

B – FMID in Base product
W – FMIDs shipped in a Web Deliverable

P – PTFs are required, P[1] – Support differs by release
P[2]- PTF required for toleration, N – Not Supported

11

© 2010 IBM Corporation

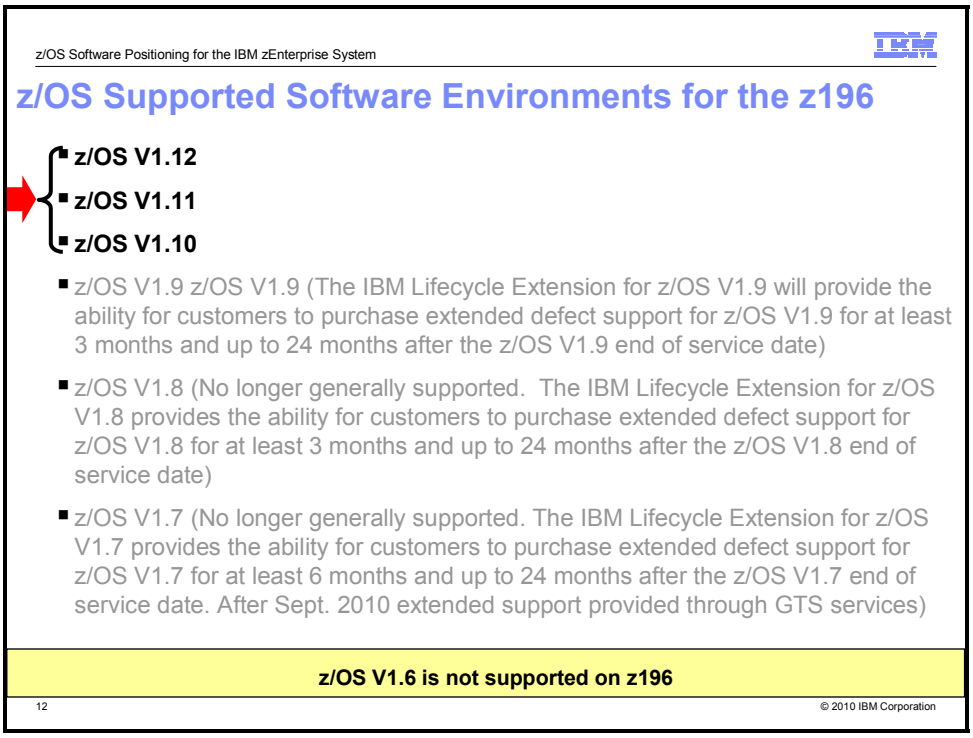

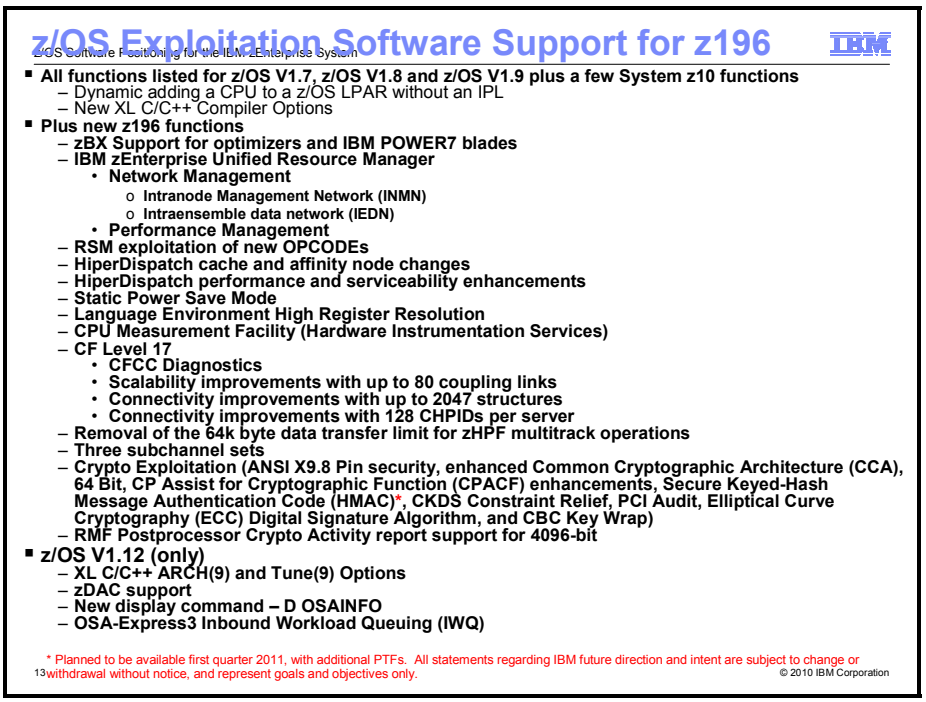

z/OS Exploitation Support for a z196 Server

z/OS V1.10 – z/OS V1.12 support is included for the following functions on the z196:

- z196 toleration support: new instruction mnemonic support, HiperDispatch, OSA-Express3 (GbE LX and SX, 1000BASE-T, 10 GbE LR and SR), FICON Express8, Parallel Sysplex InfiniBand (PSIFB) Coupling Links, High Performance FICON for System z (zHPF), Usage Report Program (IFAURP) support, toleration for >64 CPs on the server, OSA-Express3 CHPID OSD maximum ports, RMF

Postprocessor Crypto Activity report - 4096-bit, and enhancements to CPU Measurement Facility (HIS),

- IBM zEnterprise Unified Resource Manager support including: Network Management (Intranode Management Network - INMN, and Intraensemble data network - IEDN), and Performance Management,
- IBM zEnterprise BladeCenter Extension (zBX) support for IBM Smart Analytics Optimizer for DB2 for z/OS, V1.1 (5697-AQT) or IBM POWER7 Blades,
- HiperDispatch exploitation of cache and affinity node changes,
- CF Level 17 support including improved CFCC Diagnostics, Scalability improvements with up to 80 coupling links, Connectivity improvements with up to 2047 structures, and Connectivity improvements with 128 CHPIDs per server,
- Removal of the 64k byte data transfer limit for zHPF multitrack operations,
- Support for three subchannel sets,
- (z/OS V1.12 only) z/OS FICON Discovery and AutoConfiguration (zDAC) support,
- (z/OS V1.12 only) a new display command for OSAINFO,
- (z/OS V1.12 only) OSA-Express3 Inbound Workload Queuing (IWQ), support for up to 32 HiperSockets,
- (z/OS V1.12 only) XL C/C++ ARCH(9) and Tune(9) options to exploit the z196 new machine instructions
- Cryptographic exploitation if the Cryptographic Support for z/OS V1R10-V1R12 Web Deliverable (FMID HCR7780) is installed (planned availability September 2010). The cryptographic functions include: ANSI X9.8 Pin security, enhanced Common Cryptographic Architecture (CCA), 64 Bit, CP Assist for Cryptographic Function (CPACF) enhancements, Secure Keyed-Hash Message Authentication Code (HMAC) (planned to be available first quarter 2011, with additional PTFs), CKDS Constraint Relief, PCI Audit, Elliptical Curve Cryptography (ECC) Digital Signature Algorithm, and CBC Key Wrap.

z/OS Exploitation Software Support for z196

z/OS Software Positioning for the IBM zEnterprise System

	z196 PSP Bucket – 2817DEVICE 2817/ZOS																																
	Base Support																		Explicit z/OS Support															
Release	Base GA1 Support	OSA-Express3 (GbE LX and SX, 1000BASE-T, 10 GbE LR and SR)	New z/Architecture Instructions[3]	IB Coupling Links	Up to 128 Coupling Link CHPIDs	FICON Express8 (CHPID FC)	IFAURP Reporting	>64 CPs per Server	Crypto Toleration[4]	HiperDispatch	OSA-Express3 CHPID OSD maximum ports	RMF Postprocessor Crypto Activity report - 4096-bit[6]	CPU Measurement Facility (HIS)	> 64 CPs per LPAR	Cache and Affinity Node Changes	HiperDispatch Serviceability	LE High Register Resolution	Power Save Mode	CF Level 17	Extensions to zHPF Multitrack Operations	Three Subchannel Sets	Crypto Exploitation	Up to 32 HiperSockets	IBM zEnterprise BladeCenter Extension (zBX) support	Network Management	Performance Management	zDAC Support	New OSA Display Command	OSA-Express3 Inbound Workload Queuing (IWQ)	XL C/C++ ARCH(9) Tune(9)
z/OS V1.10	P	P	P[1]	P	P	P	P	P	P	B	B	P	P	P	P	P[1]	P	P	P	P	P	WP	P	P	P	P	N	N	N	N
z/OS V 1.11	P	P	P[1]	P	P	P	P	B	P	B	B	P	P	B	P	P[1]	P	P	P	P	P	WP	P	P	P	P	N	N	N	N
z/OS V 1.12	BP	BP	BP	B	BP	B	B	B	P	B	B	B	B	B	B	B	B	B	BP	B	B	WP	B	B	B	B	B	B	B	B

[1] – The Lifecycle Extension for z/OS V1.7 or z/OS V1.8 is required for support
[2] – The IBM Lifecycle Extension for z/OS V1.9 will be required after 9/30/2010
[3] – Support differs by release, z/OS V1.12 includes XL C/C++ support for ARCH(9) and TUNE(9) Options
[4] – A Crypto Web Deliverable is NOT required, but toleration PTF is needed even if a web deliverable is installed
[5] – Toleration PTFs needed to detect Power Saving Mode
[6] – PTFs recognize crypto control block, AND provide updated reports

B – FMIDin Base product
W – FMIDs shipp ed in a Web Deliverable

P – PTFs are required, P[1] – Support differs by release
P[2]- PTF required for toleration, N – Not Supported

14

© 2010 IBM Corporation

z/OS Exploitation Support for a z196 Server

IBM zEnterprise Unified Resource Manager and IBM zEnterprise BladeCenter Extension (zBX) model 002 requires at a minimum:

- z/OS V1.10 with PTFs.

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

Third subchannel set requires at a minimum:

- z/OS V1.10 with PTFs.

FICON Express8 (CHPID type FC) for support of zHPF single-track operations requires at a minimum:
- z/OS V1.10 with PTFs.

FICON Express8 (CHPID type FC) for support of zHPF multitrack operations requires at a minimum:
- z/OS V1.10 with PTFs.

FICON Express8 and FICON Express4 (CHPID type FC) supporting extension to zHPF multitrack operations requires at a minimum:
- z/OS V1.10 with PTFs.

Up to 32 HiperSockets (CHPID type IQD) requires at a minimum:
- z/OS V1.10 with PTFs.

CHPID type OSM for intranode management network requires at a minimum:
- z/OS V1.10 with PTFs.

CHPID type OSX for intraensemble data network requires at a minimum:
- z/OS V1.10 with PTFs.

12x InfiniBand and 1x InfiniBand coupling links require at a minimum:
- z/OS V1.10 with PTFs.

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 V1.10 with PTFs.

Static power save mode requires at a minimum:
- z/OS V1.10 with PTFs.

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.

Display OSAINFO for OSA-Express3 CHPID types OSD, OSM (exclusive to IBM zEnterprise 196), and OSX (exclusive to IBM zEnterprise 196) on z196 and z10 requires at a minimum:
- z/OS V1.12 with PTFs.

Inbound workload queuing (IWQ) for OSA-Express3 CHPID types OSD and OSX requires at a minimum:
- z/OS V1.12 with PTFs.

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 planned to be available September 10. 2010. This may be obtained at: http://www.ibm.com/systems/z/os/zos/downloads/

Crypto Express3 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 planned to be available September 10. 2010. This may be obtained at: http://www.ibm.com/systems/z/os/zos/downloads/

Crypto Express3 (when defined as a coprocessor) supporting **ANSI X9.8 PIN security** 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 planned to be available September 10. 2010. This may be obtained at: http://www.ibm.com/systems/z/os/zos/downloads/

Crypto Express3 (when defined as a coprocessor) supporting **Enhanced CCA key wrapping** 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 planned to be available September 10. 2010. This may be obtained at: http://www.ibm.com/systems/z/os/zos/downloads/

Crypto Express3 (when defined as a coprocessor) supporting **Secure Keyed-Hash Message Authentication Code (HMAC)** 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 planned to be available September 10. 2010. This may be obtained at: http://www.ibm.com/systems/z/os/zos/downloads/. Additional PTFs planned to be available first quarter 2011 are also required.

Crypto Express3 (when defined as a coprocessor) supporting **Elliptical Curve Cryptography (ECC) Digital Signature algorithm** requires at 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 planned to be available September 10. 2010. This may be obtained at: http://www.ibm.com/systems/z/os/zos/downloads/

© 2010 IBM Corporation

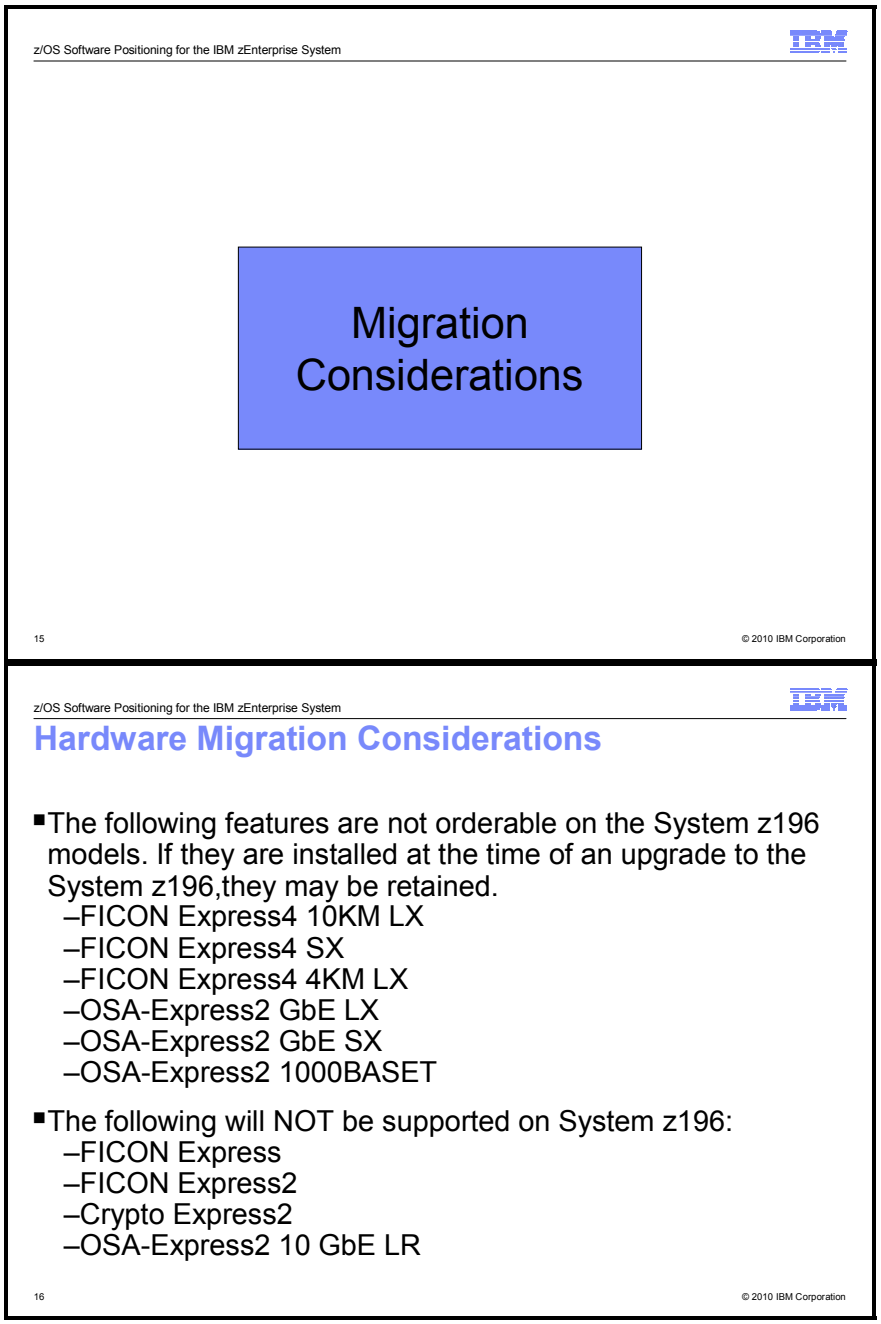

Statements of Direction

The IBM zEnterprise 196 will be the last high-end server to offer ordering of ESCON channels. This applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY. Enterprises should begin migrating from ESCON to FICON. Alternate solutions are available for connectivity to ESCON and parallel devices. IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to simplify and manage a single physical and operational environment while maximizing green-related savings.

The IBM zEnterprise 196 will be the last high-end server to support FICON Express4 channels: Enterprises should begin migrating from FICON Express4 channel features (#3321, #3322, #3324) to FICON Express8 channels.

The IBM zEnterprise 196 will be the last high-end server to support OSA-Express2 features: Enterprises should begin migrating from OSA-Express2 features (#3364, #3365, #3366) to OSA-Express3 features.

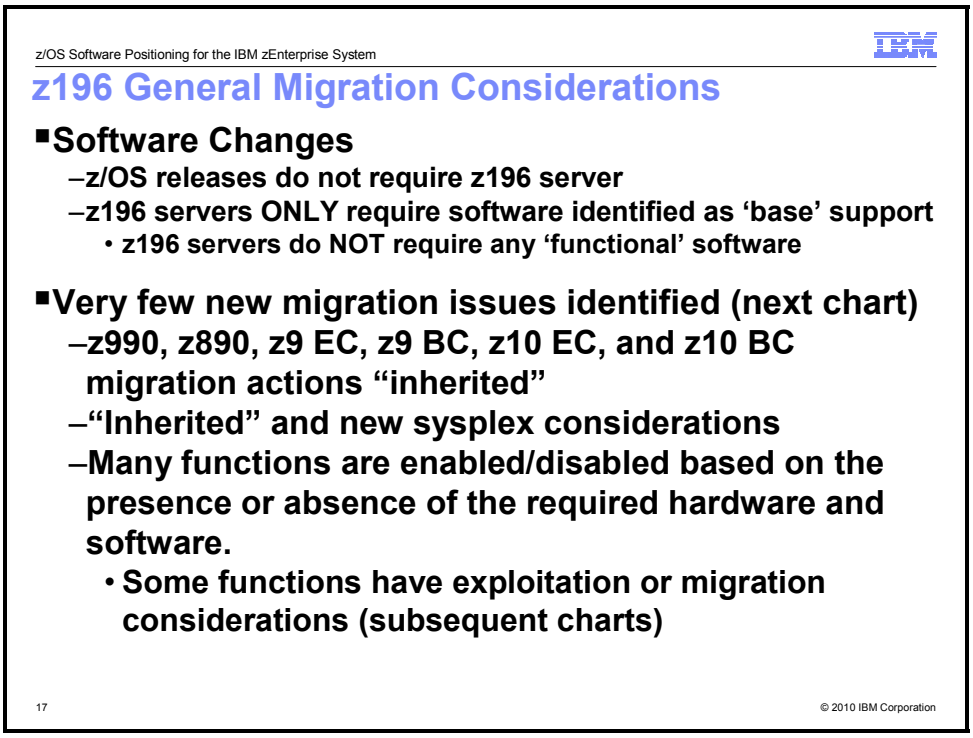

General Recommendations and Considerations
As you plan your migration to a z196 server, consider the following:

z196 z/OS Software Support
Not all System z196 functions are available in every z/OS release. Some functions have exploitation or migration considerations. Many functions are enabled or disabled, based on the presence or absence of the required hardware and software. Other than the minimum support, there is no requirement to install either software or hardware first to exploit a specific function.

Relatively few migration actions are new when coming from a z10 EC or a z10 BC server
Migration to a z196 server has, as its base, a migration to the z10 EC or z10 BC servers. This means that if you are migrating to a z196 server from a z10 EC or z10 BC server, and have performed the migration actions associated with the z10 EC or z10 BC, you have fewer migration actions than if you were migrating from a server prior to the z10 EC or z10 BC and have not yet performed the migration actions associated with these servers. There are, in fact, very few new migration actions to perform on z/OS for a z196 server if you have already migrated to a z10 EC or z10 BC server. It is important to note that you can migrate directly to a z196 server without installing the intermediate (prior to z10 EC and z10 BC) servers, but you still need to ensure that any migration considerations are satisfied for the servers that you "skipped." To read about z10 EC and z10 BC server migration actions, see *z/OS Migration to the IBM System z10* (SA23-2242-00).

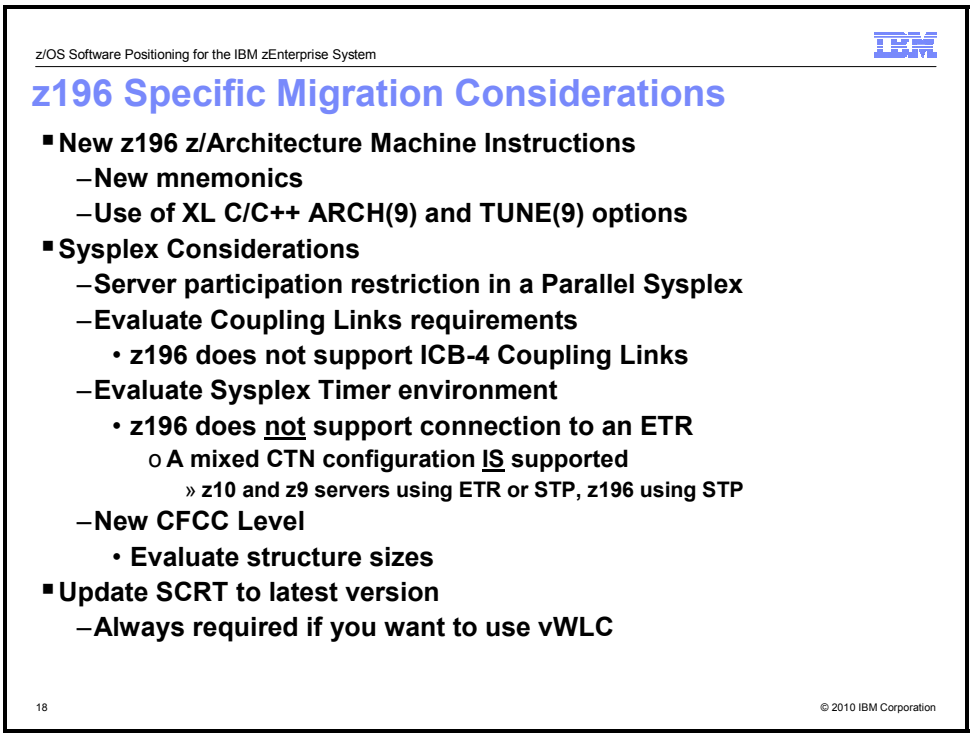

z196 Specific Migration Considerations

You can perform the following migration actions before you order or install your z196 server:

1. **Review the new mnemonics introduced for the z196 server.** The new mnemonics might collide with (be identical to) the names of assembler macro instructions you use or provide. In the event of such collisions, the HLASM's default opcode table (UNI) will treat specification of these names as instructions when APAR PK97799 is installed. This will probably cause assembler error messages and possibly cause generation of incorrect object code. See page 20 for more details.
2. **Review the sysplex configuration in which the z196 server will participate.** In particular, if you have any existing z900, z800, z990, or z890 z/OS images or coupling facility images in the sysplex, move those z/OS images or coupling facilities to later servers, such as a System z10 (z10 EC or z10 BC), System z9 (z9 EC or z9 BC), or later. This action is necessitated by the restriction that a z196 server cannot participate with prior-level servers in a sysplex.
3. **Migrate from ICB-4 to InfiniBand coupling links.** This action is necessitated because ICB-4 links are not supported on z196 servers. If desired, you can take this action after you order a z196, as you upgrade to the new server. See page 25 for more details.
4. **Implement STP (or a Mixed-CTN) timing network.** This action is necessitated because Sysplex Timers (9037-002) are not supported on a z196 server. See page 26 for more details.
5. **Migrate from unsupported hardware features to newer technology.** This action is necessitated because FICON Express, FICON Express2, Crypto Express2, and OSA-Express2 10 GbE LR are not supported on z196 servers.
6. **Run CFSIZER.** If you are moving your coupling facilities and the coupling facility structures will be on higher CFCC levels than they were previously, run the Coupling Facility Structure Sizer (CFSIZER) tool to find out if you have to increase coupling facility structure sizes. z196 servers shipped with CFCC Level 17; prepare to make the necessary changes as indicated by the tool. You can find the CFSIZER tool at http://www.ibm.com/systems/support/z/cfsizer/.
7. **Update SCRT to latest version.** If you use SCRT, make sure it is at the latest level. This is a requirement for vWLC, as well as when you upgrade servers.The latest level of SCRT can be downloaded from the SCRT Web site at http://www.ibm.com/eserver/zseries/swprice/scrt/.

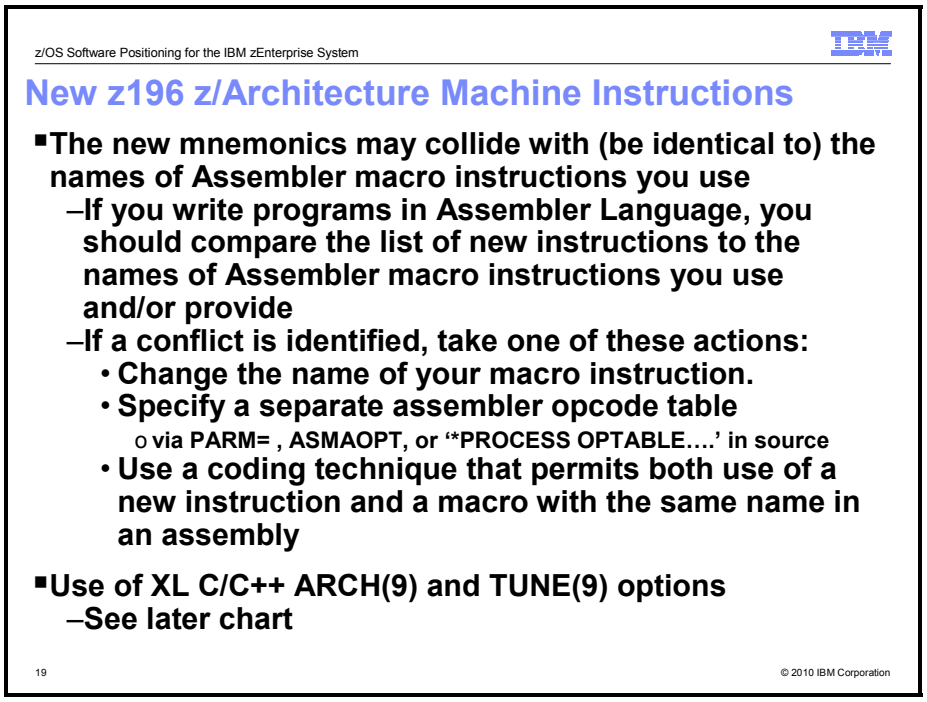

New z/Architecture Instructions

The new mnemonics might collide with (be identical to) the names of assembler macro instructions you use or provide. In the event of such collisions, the HLASM's default opcode table (UNI) will treat specification of these names as instructions when APAR PK97799 is installed. This will probably cause assembler error messages and possibly cause generation of incorrect object code.

If you write programs in Assembler Language, you should compare the list provided below and in *z/Architecture Principles of Operation* to the names of assembler macro instructions you use or provide, to identify any such conflicts or collisions that would occur following installation of HLASM APAR PK97799. If a conflict is identified, take one of the following actions:

- Change the name of your macro instruction.
- Specify PARM='...OPTABLE(YOP)...' (or some other earlier opcode table).
- Specify a separate ASMAOPT file containing assembler options, such as in the previous method (this method requires no changes to source code or JCL).
- Add, as the first statement of your source program, *PROCESS OPTABLE(YOP).
- Specify the PROFILE option either in JCL or the ASMAOPT file, and the specified or default member of the SYSLIB data set is copied into the front of the source program.
- If you must use both a new instruction and a macro with the same name in an assembly, you can use the following technique (where XXX is a sample mnemonic):

```
* Assume the default OPTABLE(UNI) is in effect
          XXX    a,b        new instruction
          PUSH   ACONTROL   save current optable definition
          ACONTROL OPTABLE(YOP)   switch optable dynamically
          XXX    r,s,t      macro invocation
          POP    ACONTROL   restore previous definition
          XXX    c,d        new instruction
```

Instruction Name	Mnemonic	Format	OPCODE
High-Word-Facility Instructions			
ADD HIGH (32)	AHHHR	RRF	B9C8
ADD HIGH (32)	AHHLR	RRF	B9D8
ADD IMMEDIATE HIGH (32)	AIH	RIL	CC8
ADD LOGICAL HIGH (32)	ALHHHR	RRF	B9CA
ADD LOGICAL HIGH (32)	ALHHLR	RRF	B9DA
ADD LOGICAL WITH SIGNED IMMEDIATE HIGH (32)	ALSIH	RIL	CCA

ADD LOGICAL WITH SIGNED IMMEDIATE HIGH (32)	ALSIHN	RIL	CCB
BRANCH RELATIVE ON COUNT HIGH (32)	BRCTH	RIL	CC6
COMPARE HIGH (32)	CHF	RXY	E3CD
COMPARE HIGH (32)	CHHR	RRE	B9CD
COMPARE HIGH (32)	CHLR	RRE	B9DD
COMPARE IMMEDIATE HIGH (32)	CIH	RIL	CCD
COMPARE LOGICAL HIGH (32)	CLHF	RXY	E3CF
COMPARE LOGICAL HIGH (32)	CLHHR	RRE	B9CF
COMPARE LOGICAL HIGH (32)	CLHLR	RRE	B9DF
COMPARE LOGICAL IMMEDIATE HIGH (32)	CLIH	RIL	CCF
LOAD BYTE HIGH (32 16)	LBH	RXY	E3C0
LOAD HALFWORD HIGH (32 16)	LHH	RXY	E3C4
LOAD HIGH (32)	LFH	RXY	E3CA
LOAD LOGICAL CHARACTER HIGH (32 16)	LLCH	RXY	E3C2
LOAD LOGICAL HALFWORD HIGH (32 16)	LLHH	RXY	E3C6
ROTATE THEN INSERT SELECTED BITS HIGH	RISBHG	RIE	EC5D
ROTATE THEN INSERT SELECTED BITS LOW	RISBLG	RIE	EC51
STORE CHARACTER HIGH (8)	STCH	RXY	E3C3
STORE HALFWORD HIGH (16)	STHH	RXY	E3C7
STORE HIGH (32)	STFH	RXY	E3CB
SUBTRACT HIGH (32)	SHHHR	RRF	B9C9
SUBTRACT HIGH (32)	SHHLR	RRF	B9D9
SUBTRACT LOGICAL HIGH (32)	SLHHHR	RRF	B9CB
SUBTRACT LOGICAL HIGH (32)	SLHHLR	RRF	B9DB
Interlocked-Access-Facility Instructions			
LOAD AND ADD (32)	LAA	RSY	EBF8
LOAD AND ADD (64)	LAAG	RSY	EBE8
LOAD AND ADD LOGICAL (32)	LAAL	RSY	EBFA
LOAD AND ADD LOGICAL (64)	LAALG	RSY	EBEA
LOAD AND AND (32)	LAN	RSY	EBF4
LOAD AND AND (64)	LANG	RSY	EBE4
LOAD AND EXCLUSIVE OR (32)	LAX	RSY	EBF7
LOAD AND EXCLUSIVE OR (64)	LAXG	RSY	EBE7
LOAD AND OR (32)	LAO	RSY	EBF6
LOAD AND OR (64)	LAOG	RSY	EBE6
LOAD PAIR DISJOINT (32)	LPD	SSF	C84
LOAD PAIR DISJOINT (64)	LPDG	SSF	C85
Load/Store-on-Condition Facility Instructions			
LOAD ON CONDITION (32)	LOC	RSY	EBF2
LOAD ON CONDITION (32)	LROC	RRF	B9F2
LOAD ON CONDITION (64)	LGOC	RSY	EBE2
LOAD ON CONDITION (64)	LGROC	RRF	B9E2
STORE ON CONDITION (32)	STOC	RSY	EBF3
STORE ON CONDITION (64)	STGOC	RSY	EBE3
Distinct-Operands-Facility Instructions			
ADD (32)	ARK	RRF	B9F8
ADD (64)	AGRK	RRF	B9E8
ADD IMMEDIATE (32 16)	AHIK	RIE	ECD8
ADD IMMEDIATE (64 16)	AGHIK	RIE	ECD9
ADD LOGICAL (32)	ALRK	RRF	B9FA
ADD LOGICAL (64)	ALGRK	RRF	B9EA
ADD LOGICAL WITH SIGNED IMMEDIATE (32 16)	ALHSIK	RIE	ECDA
ADD LOGICAL WITH SIGNED IMMEDIATE (64 16)	ALGHSIK	RIE	ECDB
AND (32)	NRK	RRF	B9F4
AND (64)	NGRK	RRF	B9E4
EXCLUSIVE OR (32)	XRK	RRF	B9F7
EXCLUSIVE OR (64)	XGRK	RRF	B9E7
OR (32)	ORK	RRF	B9F6

© 2010 IBM Corporation

OR (64)	OGRK	RRF	B9E6
SHIFT LEFT SINGLE (32)	SLAK	RSY	EBDD
SHIFT LEFT SINGLE LOGICAL (32)	SLLK	RSY	EBDF
SHIFT RIGHT SINGLE (32)	SRAK	RSY	EBDC
SHIFT RIGHT SINGLE LOGICAL (32)	SRLK	RSY	EBDE
SUBTRACT (32)	SRK	RRF	B9F9
SUBTRACT (64)	SGRK	RRF	B9E9
SUBTRACT LOGICAL (32)	SLRK	RRF	B9FB
SUBTRACT LOGICAL (64)	SLGRK	RRF	B9EB
Population-Count-Facility Instruction			
POPULATION COUNT	POPCNT	RRE	B9E1
Message Security Assist - Extension 4			
CIPHER MESSAGE WITH CFB	KMF		B92A
CIPHER MESSAGE WITH COUNTER	KMCTR		B92D
CIPHER MESSAGE WITH OFB	KMO		B92B
PERFORM CRYPTOGRAPHIC COMPUTATION	PCC		B92C
Reset-Reference-Bit-Multiple Facility			
RESET REFERENCE BIT MULTIPLE	RRBM		B9AE
Integer to/from Floating Point Convert			
CONVERT FROM FIXED (32 to extended BFP)	CXFBRA	RRF	B396
CONVERT FROM FIXED (32 to long BFP)	CDFBRA	RRF	B395
CONVERT FROM FIXED (32 to short BFP)	CEFBRA	RRF	B394
CONVERT FROM FIXED (64 to extended BFP)	CXGBRA	RRF	B3A6
CONVERT FROM FIXED (64 to long BFP)	CDGBRA	RRF	B3A5
CONVERT FROM FIXED (64 to short BFP)	CEGBRA	RRF	B3A4
CONVERT FROM LOGICAL (32 to extended BFP)	CXLFBR	RRF	B392
CONVERT FROM LOGICAL (32 to long BFP)	CDLFBR	RRF	B391
CONVERT FROM LOGICAL (32 to short BFP)	CELFBR	RRF	B390
CONVERT FROM LOGICAL (64 to extended BFP)	CXLGBR	RRF	B3A2
CONVERT FROM LOGICAL (64 to long BFP)	CDLGBR	RRF	B3A1
CONVERT FROM LOGICAL (64 to short BFP)	CELGBR	RRF	B3A0
CONVERT TO FIXED (extended BFP to 32)	CFXBRA	RRF	B39A
CONVERT TO FIXED (long BFP to 32)	CFDBRA	RRF	B399
CONVERT TO FIXED (short BFP to 32)	CFEBRA	RRF	B398
CONVERT TO FIXED (extended BFP to 64)	CGXBRA	RRF	B3AA
CONVERT TO FIXED (long BFP to 64)	CGDBRA	RRF	B3A9
CONVERT TO FIXED (short BFP to 64)	CGEBRA	RRF	B3A8
CONVERT TO LOGICAL (extended BFP to 32)	CLFXBR	RRF	B39E
CONVERT TO LOGICAL (long BFP to 32)	CLFDBR	RRF	B39D
CONVERT TO LOGICAL (short BFP to 32)	CLFEBR	RRF	B39C
CONVERT TO LOGICAL (extended BFP to 64)	CLGXBR	RRF	B3AE
CONVERT TO LOGICAL (long BFP to 64)	CLGDBR	RRF	B3AD
CONVERT TO LOGICAL (short BFP to 64)	CLGEBR	RRF	B3AC
DIVIDE TO INTEGER (short BFP)	DIEBR	RRF	B353
LOAD FP INTEGER (extended BFP)	FIXBRA	RRF	B347
LOAD FP INTEGER (long BFP)	FIDBRA	RRF	B35F
LOAD FP INTEGER (short BFP)	FIEBRA	RRF	B357
LOAD ROUNDED (extended to long BFP)	LDXBRA	RRF	B345
LOAD ROUNDED (extended to short BFP)	LEXBRA	RRF	B346
LOAD ROUNDED (long to short BFP)	LEDBRA	RRF	B344

Explanation:

- RI instruction format.
- RIE instruction format.
- RIL instruction format.
- RRE instruction format.
- RRF instruction format.
- RSY instruction format.

© 2010 IBM Corporation

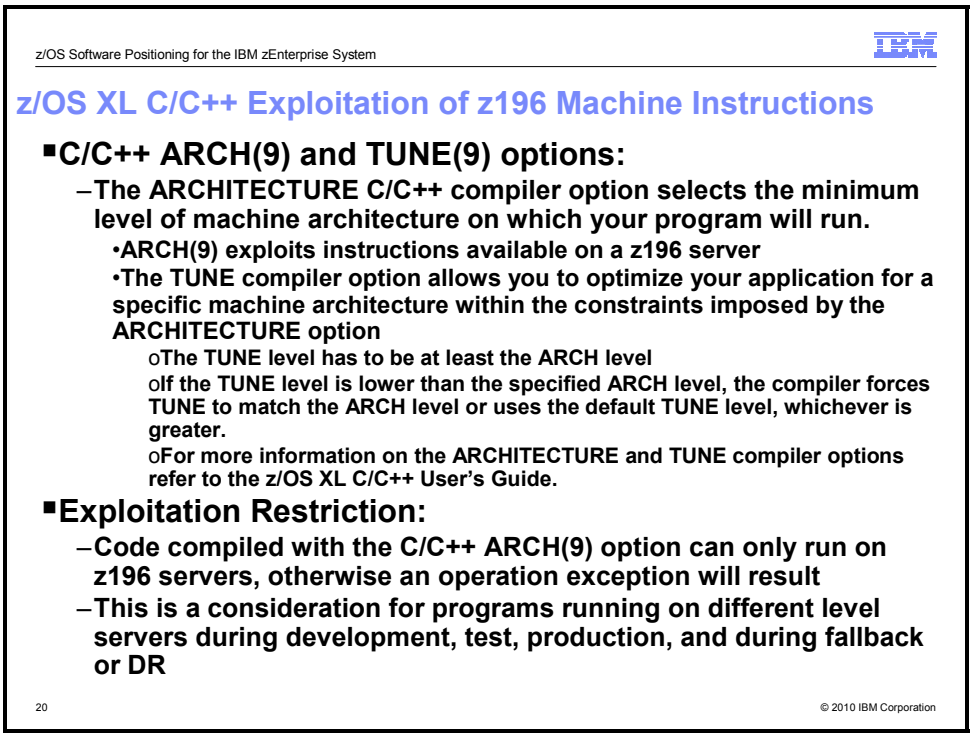

C/C++ ARCH(9) and TUNE(9) options.
The ARCHITECTURE C/C++ compiler option selects the minimum level of machine architecture on which your program can run. Certain features provided by the compiler require a minimum architecture level. ARCH(9) exploits instructions available on z196 servers. For more information, refer to the ARCHITECTURE compiler option in *z/OS XL C/C++ User's Guide*. The TUNE compiler option allows you to optimize your application for specific machine architecture. The TUNE level has to be at the ARCH level, at a minimum. If the TUNE level is lower than the specified ARCH level, the compiler forces TUNE to match the ARCH level, or uses the default TUNE level, whichever is greater. For more information about the ARCHITECTURE and TUNE compiler options refer to *z/OS XL C/C++ User's Guide*.

Exploitation restriction. When programs exploit the ARCH(9) or TUNE(9) option, those programs can only run on z196 servers, or an operation exception will occur. This is a consideration for programs that will run on different server levels (z10 and z9 servers) during development, test, and production, as well as during fallback or disaster recovery.

Server Participation in a Parallel Sysplex

- z196 does not support active participation in the same Parallel Sysplex with:
  - IBM eServer zSeries 900 (z900), IBM eServer zSeries 800 (z800),
  - IBM eServer zSeries 990 (z990), IBM eServer zSeries 890 (z890),
  - and older System/390 Parallel Enterprise Server systems

  This means:
  - Configurations with z/OS on one of these servers can't add a z196 to their sysplex for either a z/OS image or a CF image
  - Configurations with a CF on one of these servers can't add a z196 to their sysplex for either a z/OS image or a CF image
- z196 does not support ICB-4 Coupling Links
  - Customers should plan their coupling link technology
- z196 does not support connection to an ETR
  - Customers should migrate to STP ***prior*** to z196

© 2010 IBM Corporation

z196 servers in a Parallel Sysplex

z196 servers are **only** supported in a parallel sysplex with other z196 servers, z10 EC and z10 BC servers, and z9 EC and z9 BC servers. If you are running z/OS on z900, z800, z990, or z890 servers, then you cannot add a z196 server to that sysplex. That is, you will not be able to perform rolling IPLs to introduce a z196 server if you have any of the earlier (pre-System z) servers either as z/OS images or coupling facility images in the sysplex. The earlier servers in the sysplex must be upgraded to System z9 or later to have z196 supported in the sysplex. If you have any z/OS images or coupling facility images on an earlier server, and you intend to introduce a z196 into that sysplex, you must migrate those images to a System z9 (or later) server prior to introducing the z196 server.

z196 – Parallel Sysplex Coexistence Considerations

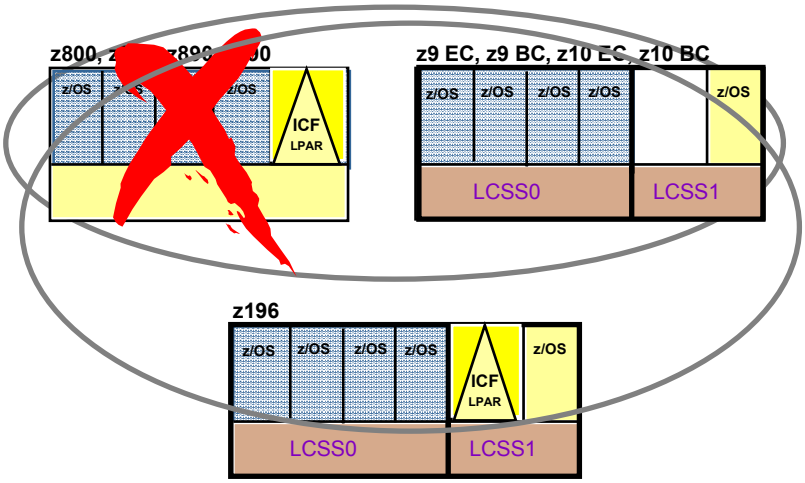

© 2010 IBM Corporation

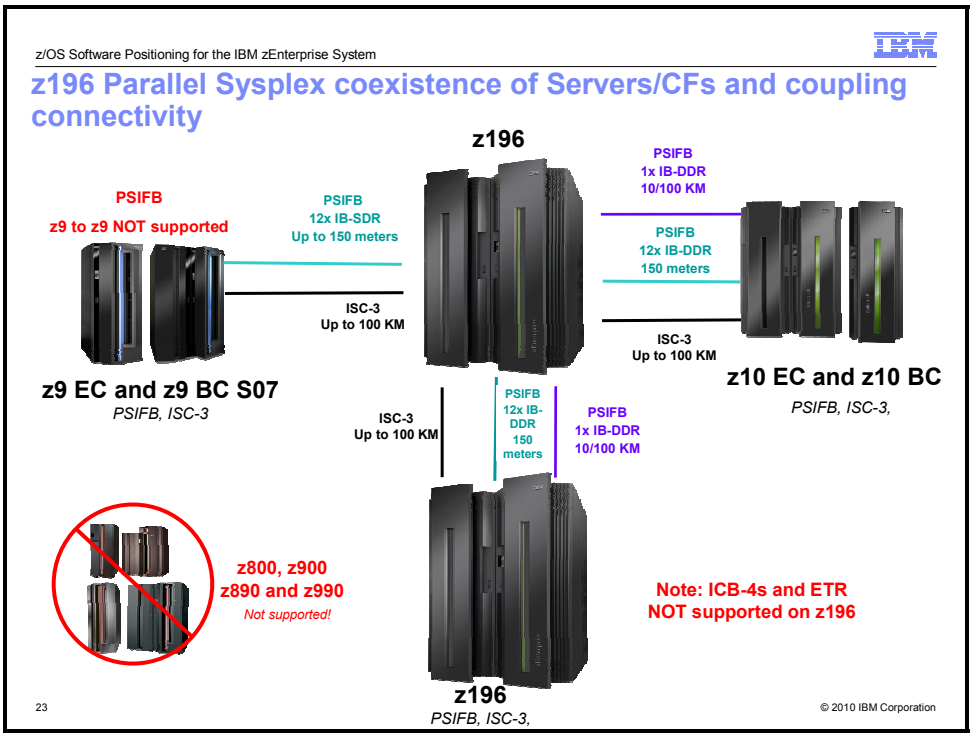

z196 does not support ICB-4 Coupling Links

The Integrated Cluster Bus 4 (ICB-4) Coupling Links are not supported on a z196 CPC. Use 12x InfiniBand coupling links, which are designed to replace Integrated Cluster Bus 4 (ICB-4), and to complement 1x InfiniBand and ISC-3 on a z196 server. InfiniBand coupling can provide significantly improved service times compared to ISC-3s for distances up to 150 meters (12x InfiniBand). You can read about InfiniBand coupling links in *IBM System z Connectivity Handbook* (SG24-5444).

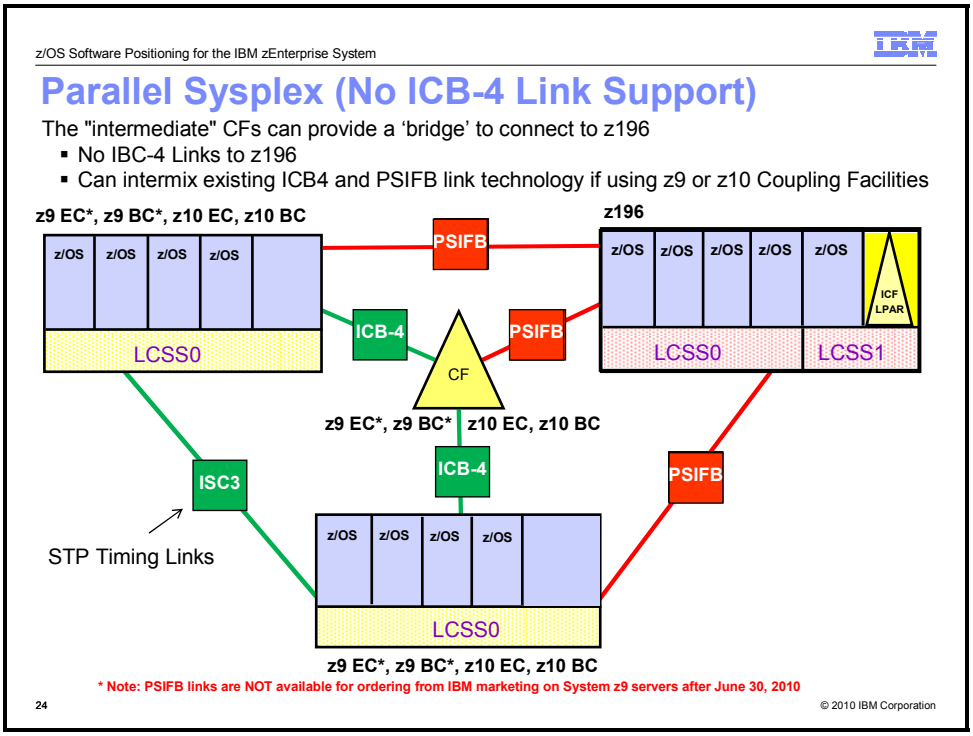

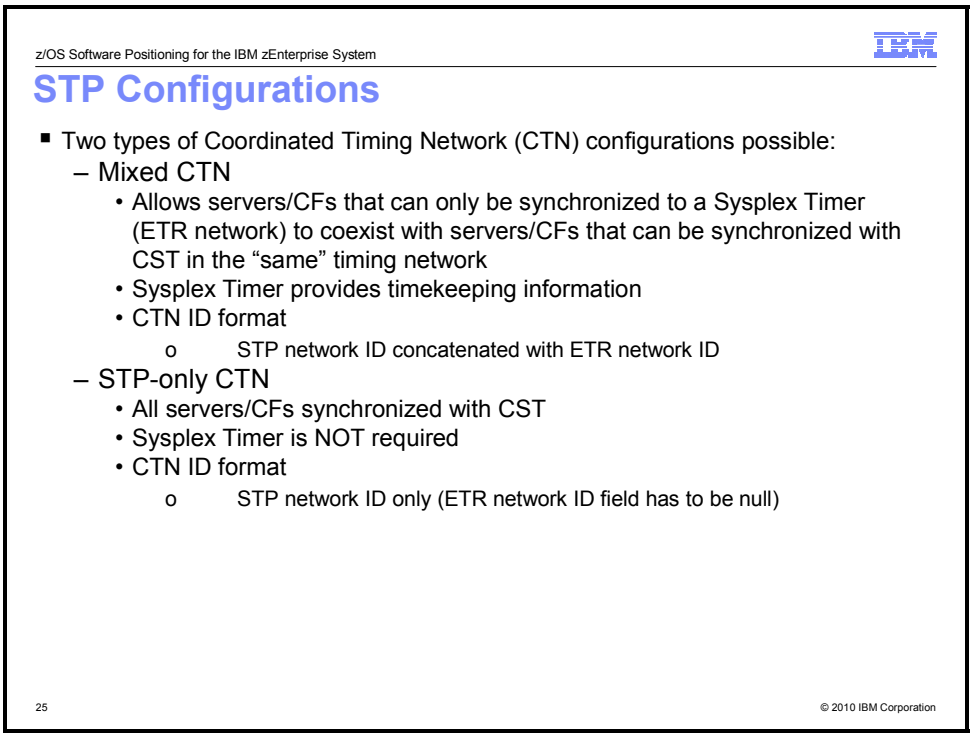

STP - time synchronization for Parallel Sysplex

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.

z196 does not support connection to an ETR

The z196 server **cannot** be connected to a Sysplex Timer (9037-002). The Server Time Protocol (STP) feature is the follow-on to the Sysplex Timer. STP is designed to allow multiple servers and coupling facilities to maintain time synchronization with each other without requiring a Sysplex Timer. STP is a hardware feature of the z196, z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890 servers. To implement STP, see the STP Web site and the publications and other resources listed there. The STP Web site is at http://www.ibm.com/systems/z/pso/stp.html.

Even though the z196 does not support attachment to a Sysplex Timer, it can participate in a Mixed-CTN that has either a z10 or z9 server synchronized to the Sysplex Timer. This maintains the capability for enterprises to concurrently migrate from an existing External Time Reference (ETR) to a Mixed-CTN and from a Mixed-CTN to an STP-only CTN.

Mixed CTN (example) with System z196 server

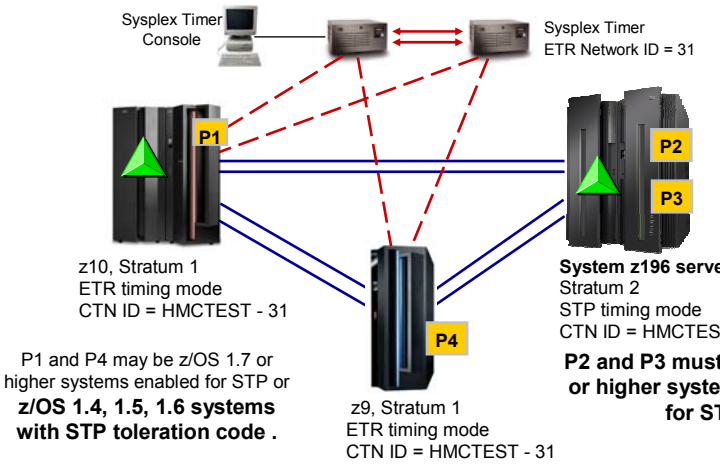

Strong recommendation: Configure at least two Stratum 1 servers before configuring a z196 server as S2 – to avoid a single point of failure

© 2010 IBM Corporation

STP-only CTN (example) with System z196 server

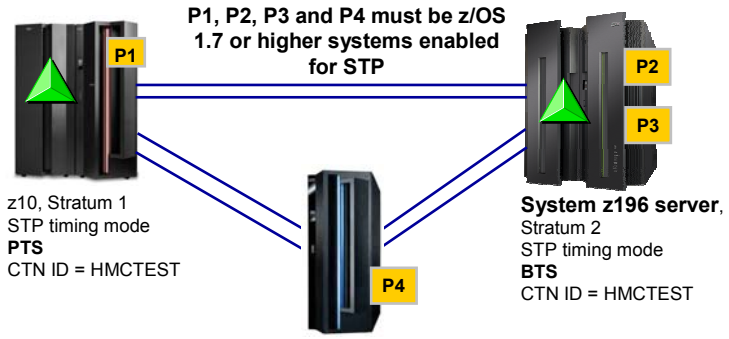

© 2010 IBM Corporation

System z CFCC Level 17

- **CFCC Level 17 allows:**
  - Up to 2047 CF structures (CFCC 16 allowed 1023).
    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
  - Improved CFCC Diagnostics & Link Diagnostics
- **Structure and CF Storage Sizing with CFCC level 17**
  - May increase storage requirements when moving from:
    - CF Level 16 (or below) to CF Level 17
  - CF Sizer Tool recommended
    - http://www.ibm.com/systems/z/cfsizer/
- **Greater than 1023 CF Structures requires a new version of the CFRM CDS**
  - All systems in the sysplex need to be at z/OS V1.12 or have the coexistence/preconditioning PTF installed.
  - Falling back to a previous level (without coexistence PTF installed) is NOT supported without sysplex IPL

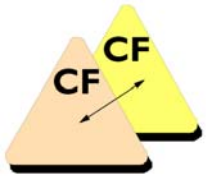

28

© 2010 IBM Corporation

Parallel Sysplex Enhancements

Connectivity improvements with up to 80 coupling links: The IBM zEnterprise 196 increases the number of external coupling links allowed from 64 to 80. This allows the full configuration of 32 PSIFB 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 z196, we have 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 (CFCC) 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. z196 and CFCC level 17 provide improved serviceability of Coupling Facilities with enhanced data collection and triggering of nondisruptive CF dumps.

Larger coupling facility structure sizes might be necessary. When you change coupling facility control code (CFCC) levels, your coupling facility structure sizes might change. z196 servers initially ship with CFCC level 17. If, as part of your migration to a z196 server, you change CFCC levels (either by placing a coupling facility on the z196 server or by moving the coupling facility to a z10 EC or z10 BC at a later CFCC level), you might have larger structure sizes than you did previously. In this case you should run the Coupling Facility Structure Sizer (CFSIZER) tool to find out if you have to increase coupling facility structure sizes. Prepare to make the necessary changes as indicated by the tool. You can find the CFSIZER tool at http://www.ibm.com/systems/support/z/cfsizer/.
Note: The PTFs to support CFCC Level 17 have coexistence (or sysplex preconditioning) PTFs that require installation throughout your sysplex prior to implementing CFCC Level 17.

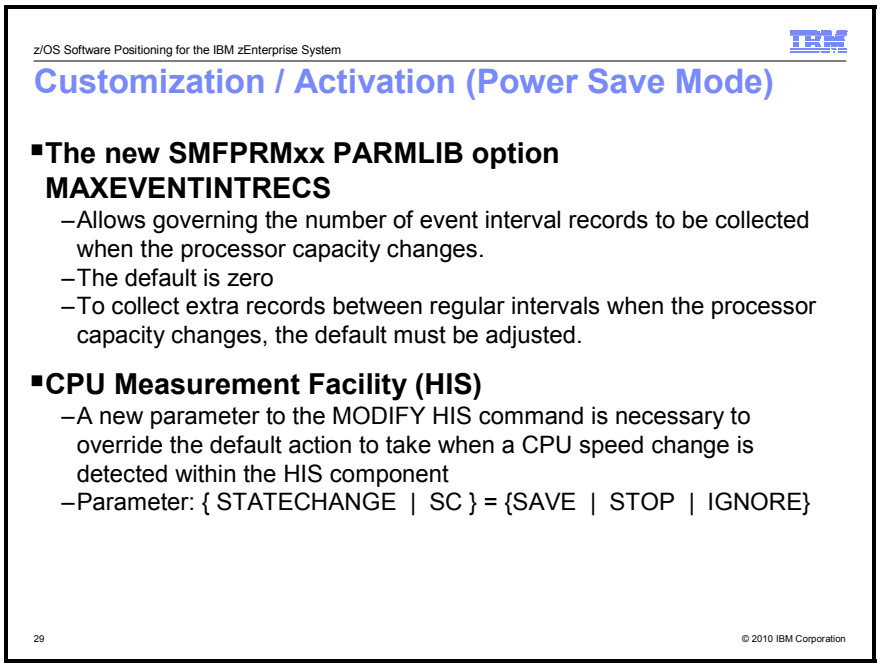

Static Power Save Mode

With the introduction of static power save mode, IBM is continuing to help provide energy management capabilities to facilitate energy savings. You now have the ability to reduce the power consumption of the IBM zEnterprise 196.

Static power save mode is designed to reduce power consumption on z196 when full performance is not required. It can be switched on and off during runtime with no disruption to currently running workloads, aside from the change in performance. On air-cooled models, Static power save mode can be entered once in a calendar day.

Using the Hardware Management Console (HMC), as well as the Active Energy Manager (AEM), you can use static power save mode for:

- Periods of lower utilization - weekends, third shift
- Capacity backup systems - systems used for emergency backup; keep them "running" but reduce energy consumption. Systems can quickly be brought back to full performance.

This could result in a 20% - 30% reduction in power consumption (depending on system configuration). If you have also implemented dynamic control of cooling (for example, through the integration of Active Energy Manager with facility management applications) the air conditioning for the cooling zone can be turned down and additional energy can be saved.

There is a new SMFPRM*xx* parmlib option, MAXEVENTINTRECS, that allows governing the number of event interval records to be collected when the processor capacity changes. The default is zero. To collect extra records between regular intervals when the processor capacity changes, the default must be adjusted.

If you are using the CPU Measurement Facility (Hardware Instrumentation Services), there is a new parameter on the MODIFY *hisproc* command. You can use this parameter, STATECHANGE, to override the default action to take when a CPU speed change is detected within the HIS component.

- STATECHANGE is an optional parameter specifying the action to take when a significant change to the system is detected. An example of a significant state change would be if the capability of a CPU has changed. When a state change is detected, the following actions are allowed:
  - **STOP** - When you specify STOP, the current run of hardware event data collection for a system is ended as if the DURATION of the run has expired.
  - **IGNORE** - When you specify IGNORE, the current run of hardware event data collection continues as if no state change had been detected.
  - **SAVE** - When you specify SAVE, the current run of hardware event data collection will first save all data collected as if the DURATION of the run has expired, then it will resume collecting hardware event data in the new state for the remaining DURATION of the current collection run. (Default).

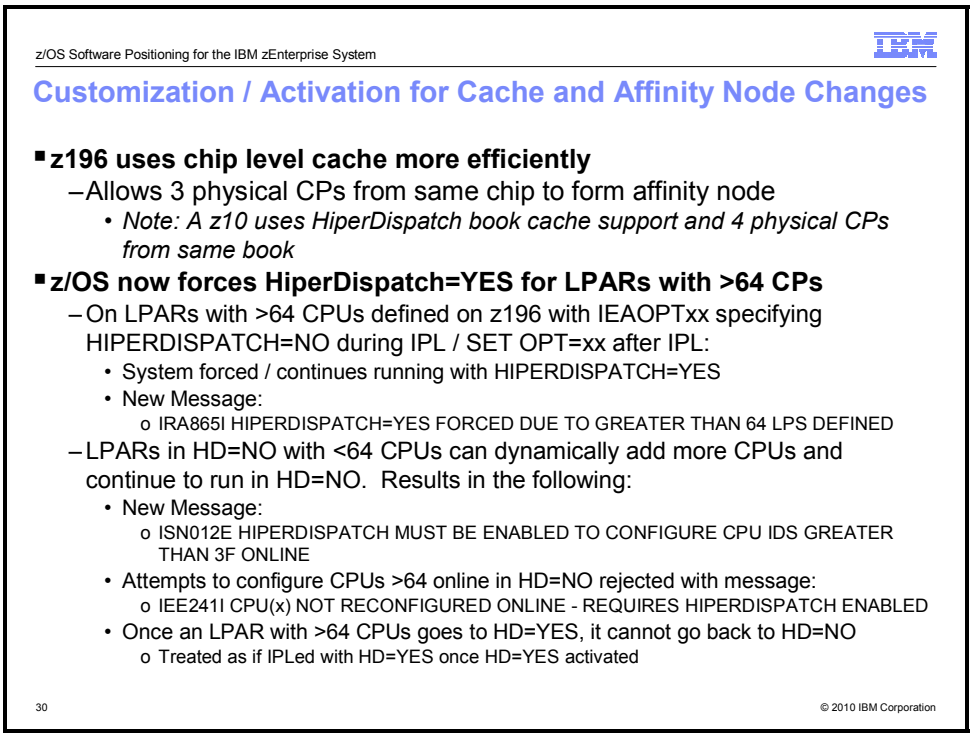

HiperDispatch cache and affinity node changes
This function is enhanced to exploit the z196 architecture and now allows three physical CPs from same chip to form affinity node. A z10 uses HiperDispatch book cache support and four physical CPs from same book.

To realize the benefits of HiperDispatch, z/OS has been changed to force HiperDispatch=YES for LPARs with greater than 64 CPUs. On LPARs with greater than 64 CPUs defined on a z196 server with IEAOPT*xx* specifying HIPERDISPATCH=NO during IPL (or SET OPT=*xx* after IPL), the system generates a message but continues to run with HIPERDISPATCH=YES. The new message is IRA865I HIPERDISPATCH=YES FORCED DUE TO GREATER THAN 64 LPS DEFINED.

On LPARs in which HIPERDISPATCH=NO is specified with less than 64 CPUs, you can dynamically add more CPUs and continue to run in HIPERDISPATCH=NO. However, you may see the new message ISN012E HIPERDISPATCH MUST BE ENABLED TO CONFIGURE CPU IDS GREATER THAN 3F ONLINE.

Any attempt to configure CPUs greater than 64 CPUs online in HIPERDISPATCH=NO will be rejected with message IEE241I CPU(x) NOT RECONFIGURED ONLINE - REQUIRES HIPERDISPATCH ENABLED.

An LPAR with greater than 64 CPUs that dynamically changed to HIPERDISPATCH=YES cannot go back to HIPERDISPATCH=NO. It will be treated as if it was IPLed with HIPERDISPATCH=YES after HIPERDISPATCH=YES is activated.

To assist with warning when you are getting close to 64 CPUs and running with HIPERDISPATCH=NO, the IBM Health Checker for z/OS check, SUP_HiperDispatchCPUConfig, is added in z/OS V1R12 and available on z/OS V1R11 and z/OS V1R10 with APAR OA30476. The check always succeeds for LPAR in HIPERDISPATCH=YES (all CPU configurations supported). When an LPAR is running with HIPERDISPATCH=NO, the check raises an exception when the number of CPUs is close to forcing the LPAR to IPL with HIPERDISPATCH=YES. The CPUSLEFTB4NEEDHD parameter indicates the minimum number of CPUs that can be installed and activated on an LPAR running in HIPERDISPATCH=NO. When CPUSLEFTB4NEEDHD=0, the check always succeeds. The default is 8, with values 0-63 accepted. The system redrives the check when the HIPERDISPATCH state changes or CPUs are dynamically added. Possible IBM Health Checker for z/OS messages:

- IEAVEH080I CPU configuration supported with HiperDispatch *curstate*
- IEAVEH081E CPU configuration supported with HiperDispatch disabled.

numcpus more CPU(s) can be added with HiperDispatch disabled.

Customization / Activation for Cache and Affinity Node Changes

- **On z196, SUP_HiperDispatchCPUConfig Health Check added**
  - Always succeeds for LPAR in HD=YES (all CPU configurations supported)
  - When LPAR in HD=NO, raises an exception when # of CPUs is "close" to forcing the LPAR to IPL with HD=YES.
    - CPUSLEFTB4NEEDHD (default 8, 0-63 accepted) parameter indicates the minimum number of CPUs that can be installed and activated on an LPAR running in HD=NO.
      - When CPUSLEFTB4NEEDHD=0, the check always succeeds.
  - System redrives check when HD state changes or CPU(s) dynamically added
- **Healthchecker Messages:**
  - IEAVEH080I CPU configuration supported with HiperDispatch *curstate*
  - IEAVEH081E CPU configuration supported with HiperDispatch disabled. *numcpus* more CPU(s) can be added with HiperDispatch disabled.

© 2010 IBM Corporation

Customization / Activation for HiperDispatch Serviceability

- **Provides HiperDispatch knobs, SIGP data, and suspend lock data needed to debug performance problems**
  - **z/OS Version 1 Release 11**
    - System suspend lock instrumentation data collection
    - Count SIGPs contributing to LPAR overheads data collection
  - **z/OS Version 1 Release 12**
    - System latch contention instrumentation data collection
    - Data collection for RMF for WEB Queue Distribution
- **There are no actions required to enable most of the functions**
  - Data is always collected
  - RMF can be configured to report some of the data

© 2010 IBM Corporation

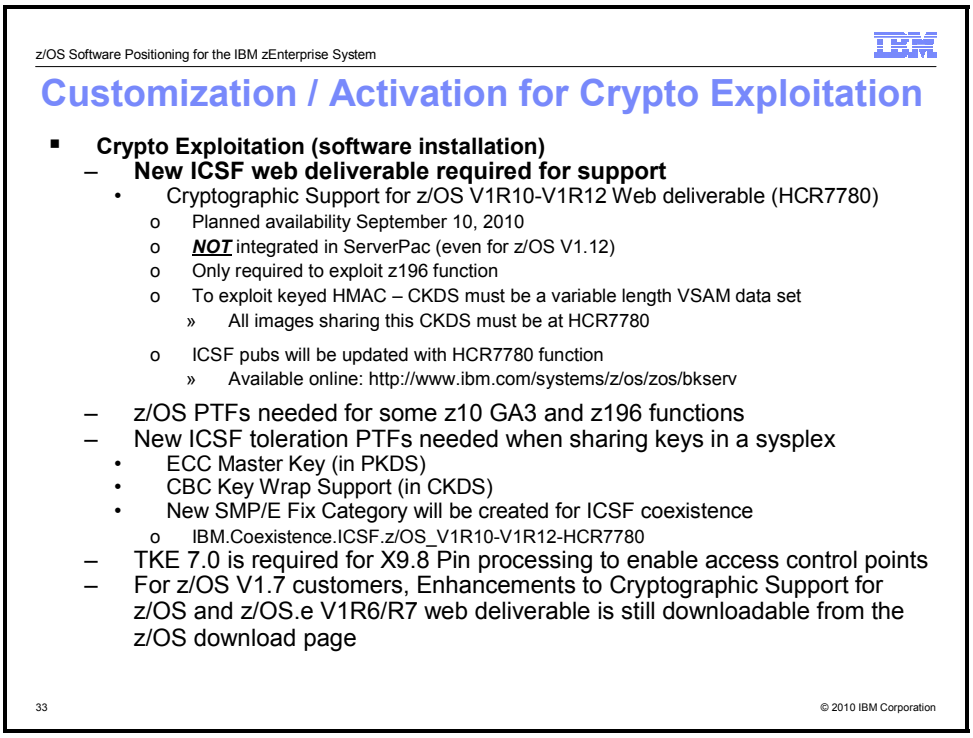

CP Assist for Cryptographic Function (CPACF) enhancements:
The following are exploitation of Message-Security-Assist Extension 4:
New instructions:

1. Cipher Message with CFB (KMF)
2. Cipher Message with Counter (KMCTR)
3. Cipher Message with OFB (KMO)

New function codes for existing instructions:

- Compute intermediate Message Digest (KIMD) adds KIMD, an 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, (RFA51029) dated October 20, 2009.

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 z196, 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, (RFA51029) dated October 20, 2009.

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 z196 and is supported by z/OS and z/VM.
- **Enhanced Common Cryptographic Architecture (CCA) - a key wrapping to comply with ANSI X9.24-1 key bundling requirements**, 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 z196 and is supported by z/OS and z/VM.

- **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 z196 and is supported by z/OS and z/VM. For z/OS support is to be made available with the Cryptographic Support for z/OS V1.10 through z/OS V1.12 Web deliverable with the additional PTF for APAR OA33260 which is planned for February 2011 availability.
- **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 z196 and is supported by z/OS and z/VM.

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 z196 and is supported by z/OS and z/VM.

Trusted Key Entry 7.0 Licensed Internal Code (LIC)

- ECC Master Key Support: From the Trusted Key Entry (TKE) workstation, administrators can generate key material, load or clear the new ECC master key register, or clear the old ECC master key register. The ECC key material can be stored on the TKE or on a smart card.
- CBC Default Settings Support: The TKE provides function that allows the TKE user to set the default key wrapping method used by the host crypto module.
- TKE Audit Record Upload Configuration Utility Support: The TKE Audit Record Upload Configuration Utility allows TKE workstation audit records to be sent to a System z host and saved on the host as z/OS System Management Facilities (SMF) records. The SMF records have a record type of 82 (ICSF) and a subtype of 29. TKE workstation audit records are sent to the same TKE Host Transaction Program that is used for Trusted Key Entry operations.
- USB Flash Memory Drive Support: The TKE workstation now supports a USB flash memory drive as a removable media device. When a TKE application displays media choices, the application allows you to choose a USB flash memory drive if the IBM-supported drive is plugged into a USB port on the TKE and it has been formatted for the specified operation.
- Stronger PIN Strength Support: TKE smart cards created on TKE 7.0 require a 6-digit PIN rather than a 4-digit PIN. TKE smart cards that were created prior to TKE 7.0 continue to use 4-digit PINs and will work on TKE 7.0 without changes. You can take advantage of the stronger PIN strength by initializing new TKE smart cards and copying the data from the old TKE smart cards to the new TKE smart cards.
- Stronger Password Requirements for TKE Passphrase User Profile Support: New rules are required for the passphrase used for passphrase logon to the TKE workstation crypto adapter. The passphrase must:
  - Be 8 to 64 characters long
  - Contain at least 2 numeric and 2 non-numeric characters
  - Not contain the user ID

  These rules are enforced when you define a new user profile for passphrase logon, or when you change the passphrase for an existing profile. Your current passphrases will continue to work.

© 2010 IBM Corporation

z196 Crypto Express3 – UDX

- **UDX (User Defined eXtension)**
  - –Extends the functionality of IBM's CCA (Common Cryptographic Architecture) application program
    - Customized cryptographic verb controls per customer
  - –UDX interfaces using HW control blocks and ICSF CB, therefore if hw platform changes or ISCF level changes or both, then UDX must updated for the new control blocks
    - If a customer has UDX, they would already know this

© 2010 IBM Corporation

© 2010 IBM Corporation

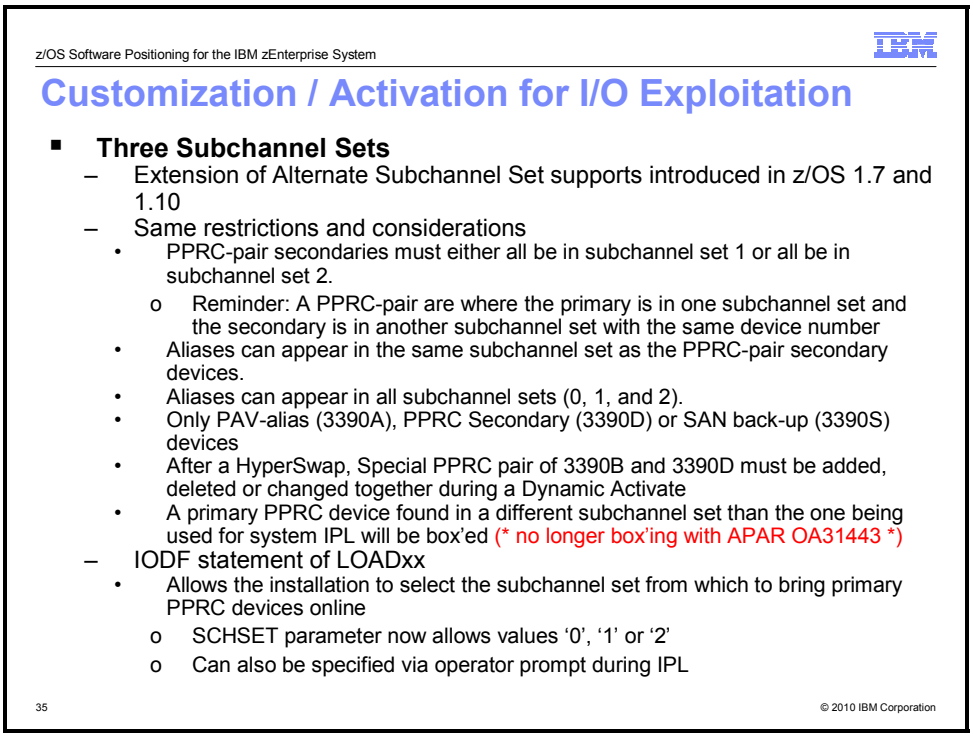

Third subchannel set

You now have the ability to extend the amount of addressable storage capacity to help facilitate storage growth with the introduction of a third subchannel set - an additional 64K subchannels - to help complement other functions such as "large" or extended addressing volumes and HyperPAV. This may also help to facilitate consistent device address definitions, simplifying addressing schemes for congruous devices.

The first subchannel set (SS0) allows definitions of any type of device (such as bases, aliases, secondaries, and those other than disk that do not implement the concept of associated aliases or secondaries). The second and third subchannel sets (SS1 and SS2) can now both be designated for use for disk alias devices (of both primary and secondary devices) and/or Metro Mirror secondary devices only.

The third subchannel set applies ESCON, FICON and zHPF protocols, and is supported by z/OS and Linux on System z.

Definitions for the third subchannel set are similar to those for the second subchannel set and can be made with HCD.

The IODF statement of LOAD*xx* allows the installation to select the subchannel set from which to bring primary PPRC devices online. The SCHSET parameter now allows values **0, 1**, or **2**. It can also be specified through an operator prompt during IPL.

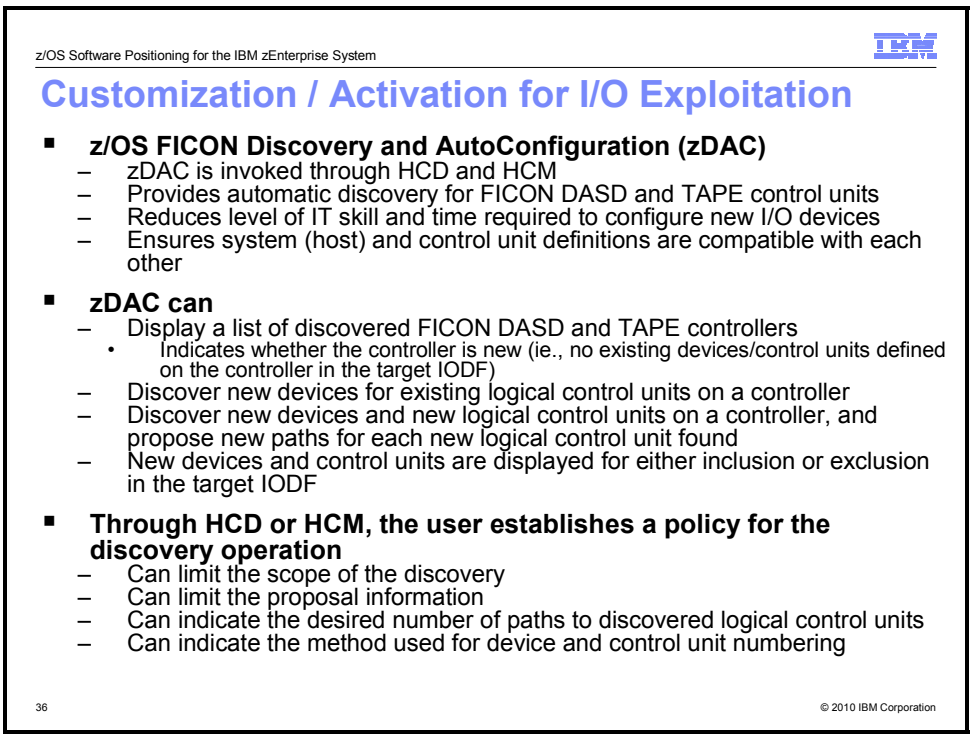

z/OS Discovery and AutoConfiguration (zDAC) for FICON Channels:
A new IBM System z Discovery and AutoConfiguration (zDAC) function, available with z/OS V1.12 on IBM zEnterprise 196 server, can help simplify the configuration for new and changed disk and tape I/O configuration changes. zDAC can save time by discovering new and changed fabric-attached devices, and suggesting configurations aligned with best practices for availability and with the I/O policies that you set. For example, with zDAC, adding new controllers to an existing I/O configuration can take as little as a few minutes. 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 a single system or for all the systems in a sysplex that support the function and does not alter your existing configuration until you have accepted or modified the proposed changes. zDAC proposes configurations that incorporate not only the current contents of your work I/O definition file (IODF) with additions for new and changed controllers and devices, but also IBM best practices for I/O high availability. zDAC is designed to help simplify I/O configuration on z196 running z/OS and reduce complexity and setup time.

zDAC applies to all FICON features supported on z196 when configured as CHPID type FC and is supported by z/OS. zDAC is exploited by the IBM System Storage DS8700 which is designed to improve discovery performance. zDAC is supported by all FICON-attached IBM System Storage tape and disk control units.

To use zDAC, you must first establish a policy for the discovery operation. This is done through HCD or HCM. You can limit the scope of the discovery, limit the proposal information, indicate the desired number of paths to discovered logical control units, and indicate the method used for device and control unit numbering. After controllers and devices are discovered, you can select which controllers to be defined and accept or override the proposed values for control units and devices.

Customization / Activation for I/O Exploitation

- **z/OS FICON Discovery and AutoConfiguration (zDAC)**
  - zDAC is defined using the HCD Panels.
  - In order to run zDAC, each TSO userid has to have a HCD Profile set up.
  - The HCD profile is accessed HCD Profile Options and Policies (Option 0)
    - The Autoconfiguration Policy is Option 0,2 (next chart)
    - To add or edit Autoconfiguration Logical Partition Groups is Option 0,3
      - The LPARs specified in the LP Group can span z196 CECs
      - **All members of the group must be running z/OS V1.12 or else the function will not work.**
    - To add or edit Autoconfiguration Operating System Groups is Option 0,4
      - Multiple OS's can be targeted for zDAC discovery.

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration Policy

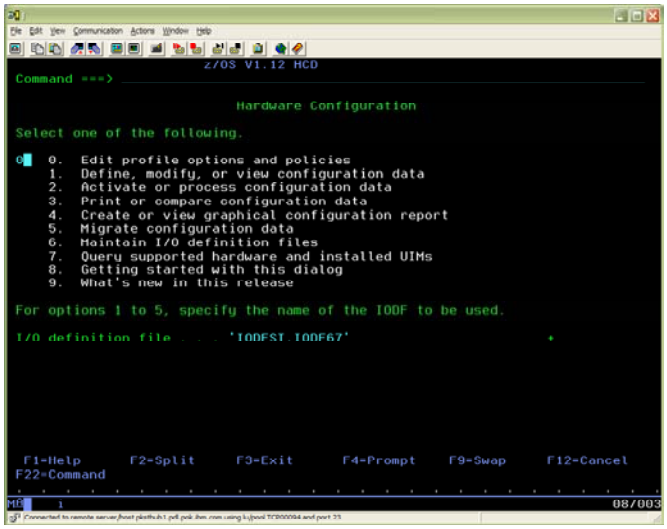

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration Policy

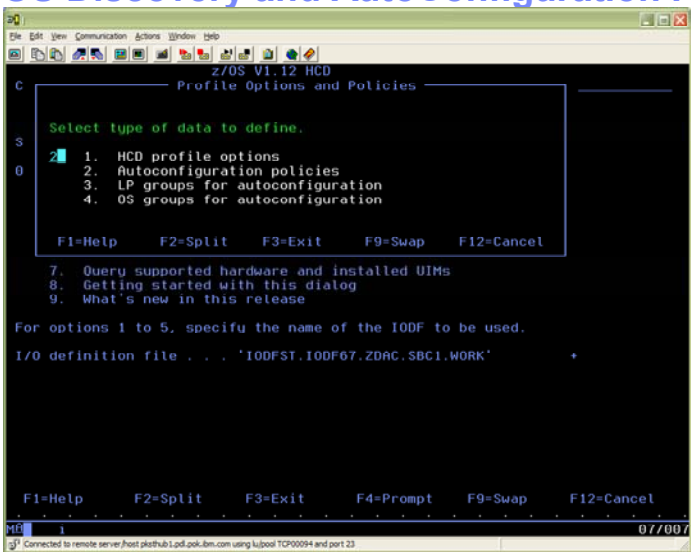

© 2010 IBM Corporation

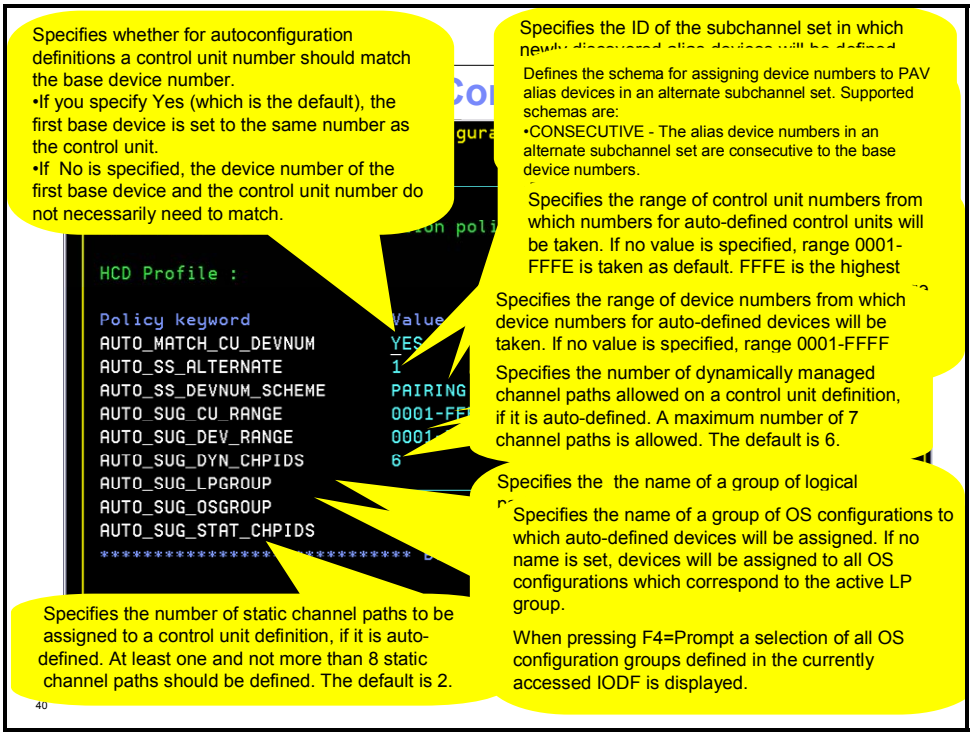

AUTO_MATCH_CU_DEVNUM - Specifies whether for autoconfiguration definitions a control unit number should match the base device number.

- If you specify Yes (which is the default), the first base device is set to the same number as the control unit.
- If No is specified, the device number of the first base device and the control unit number do not necessarily need to match.

AUTO_SS_ALTERNATE - Specifies the ID of the subchannel set in which newly discovered alias devices will be defined during an auto-definition process, provided that free device numbers are available in this subchannel set, and processors that have access to the device range, support alternate subchannels. The default subchannel set ID is 1.
AUTO_SS_DEVNUM_SCHEME - Defines the schema for assigning device numbers to PAV alias devices in an alternate subchannel set. Supported schemas are:

- CONSECUTIVE - The alias device numbers in an alternate subchannel set are consecutive to the base device numbers.
- DENSE - The device numbers in an alternate subchannel set are densely assigned, that is the next free device numbers in the assigned device number range will be used.
- PAIRING - Base and alias device numbers are assigned alternatively starting with for example device numbers xx00 and xx80 versus xx80 and xx00. PAIRING is the default.

AUTO_SUG_CU_RANGE - Specifies the range of device numbers from which device numbers for auto-defined devices will be taken. If no value is specified, range 0001-FFFF is taken as default.
AUTO_SUG_DEV_RANGE - Specifies the range of control unit numbers from which numbers for auto-defined control units will be taken. If no value is specified, range 0001-FFFE is taken as default. FFFE is the highest supported control unit number. Specify the range according to the following syntax: nnnn-mmmm where nnnn is the lower range boundary, and mmmm is the upper range boundary.
AUTO_SUG_DYN_CHPIDS - Specifies the number of dynamically managed channel paths allowed on a control unit definition, if it is auto-defined. A maximum number of 7 channel paths is allowed. The default is 6.
AUTO_SUG_LPGROUP - Specifies the name of a group of logical partitions to which auto-defined devices will be assigned. If no name is set, devices will be assigned to all partitions of the active sysplex. Reserved group name ALL signals this during autoconfiguration processing. When pressing F4=Prompt a selection of all logical partition groups defined in the currently accessed IODF is displayed.
AUTO_SUG_OSGROUP - Specifies the name of a group of OS configurations to which auto-defined devices will be assigned. If no name is set, devices will be assigned to all OS configurations which correspond to the active LP group. When pressing F4=Prompt a selection of all OS configuration groups defined in the currently accessed IODF is displayed.
AUTO_SUG_STAT_CHPIDS - Specifies the number of static channel paths to be assigned to a control unit definition, if it is auto-defined. At least one and not more than 8 static channel paths should be defined. The default is 2.

© 2010 IBM Corporation

zDAC – HCD Add AutoConfiguration LP Group

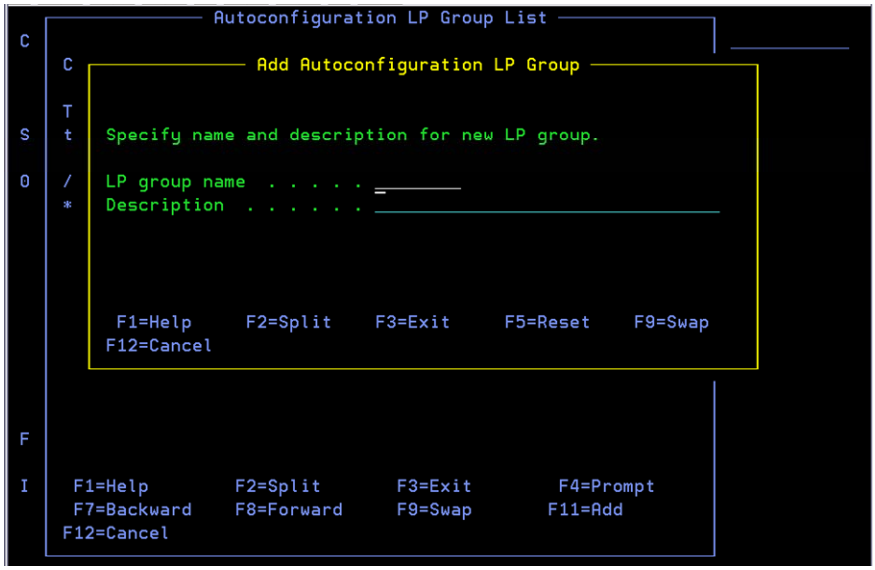

© 2010 IBM Corporation

zDAC – HCD Add AutoConfiguration OS Group

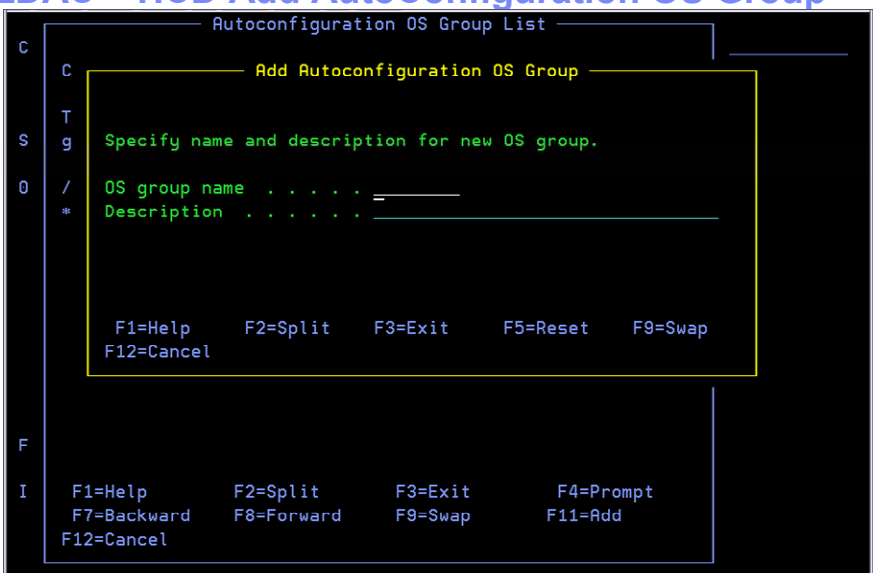

© 2010 IBM Corporation

Customization / Activation for I/O Exploitation

zDAC

- zDAC discovery process is invoked via HCD panel Options 1,6.
- There are two phases of zDAC discovery:
  1. Fabric Discovery Phase
     - This phase interrogates the switches (MdData, Brocade and CISCO) and discovers all of the TAPE and DASD Nodes that are logged in to the fabric.
     - Once this step is executed, a list of nodes is presented for the next phase of discovery.
  2. Controller Discovery Phase
     - This phase interrogates the Node, Tape or DASD, when selected for discovery. It detemines CU and Device numbering based on the CU configuration, and the paths to be used for access.
     - A proposal will be presented and when accepted, the IODF workfile will contain the entries for the newly discovered node(s).
- Nodes that are discovered during the Fabric Discovery Phase are characterized as either NEW=YES or NEW=NO.
  - NEW=YES means that none of the devices on that Node have been defined in the IODF used for Discovery.
  - NEW=NO means that one or more devices on that Node are currently defined in the IODF used for discovery.

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration Request

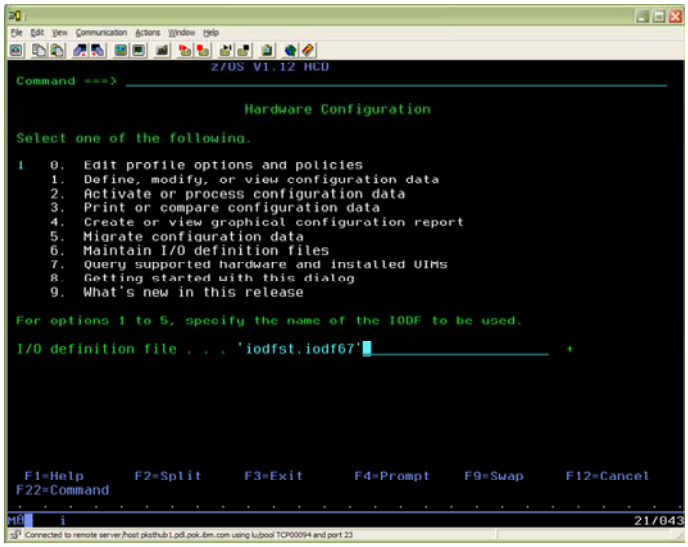

© 2010 IBM Corporation

zDAC – Discovery Process

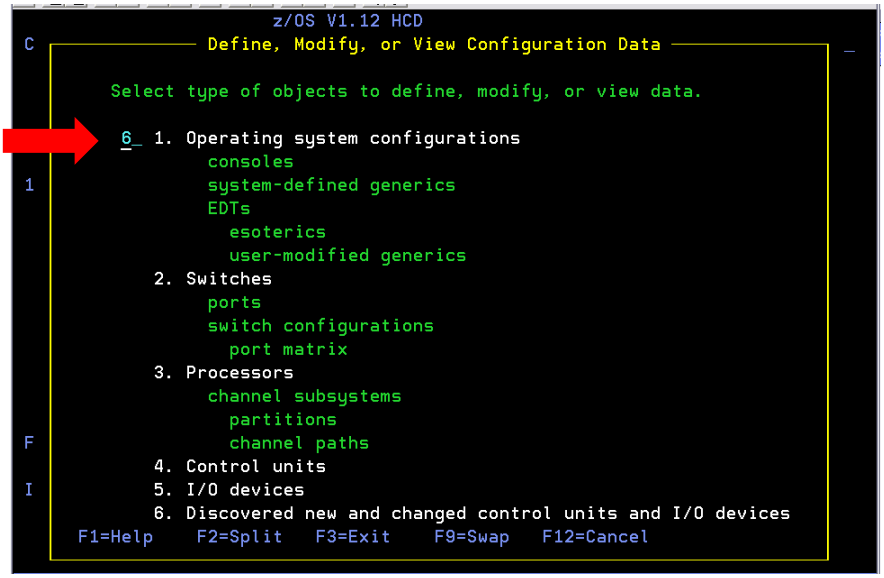

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration Request

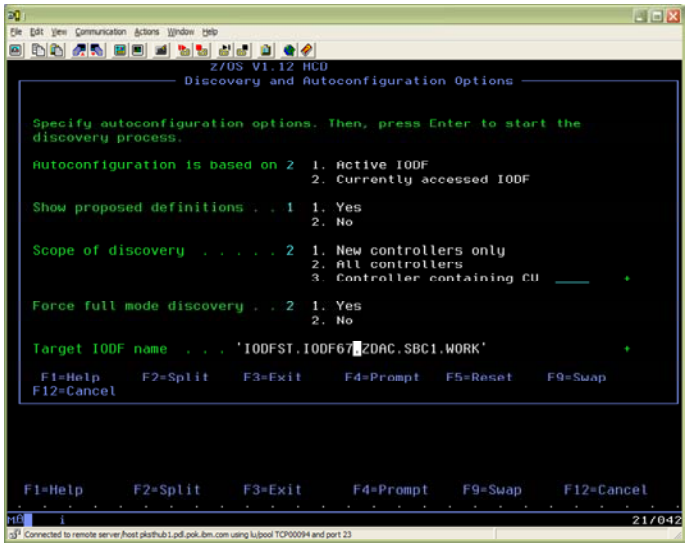

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration

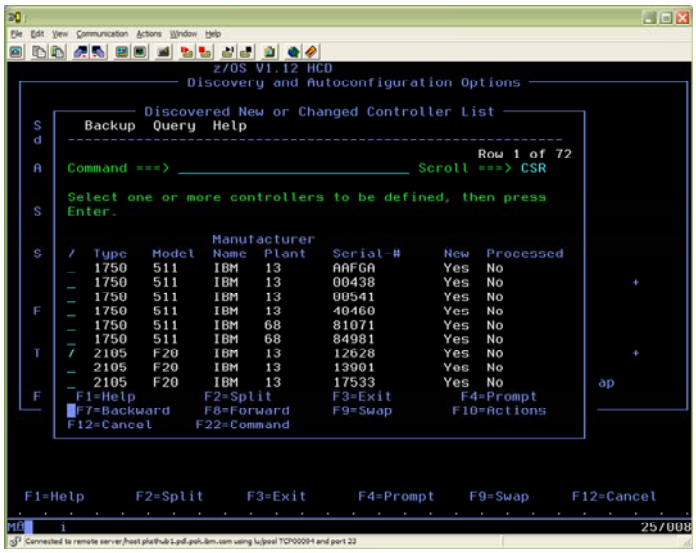

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration

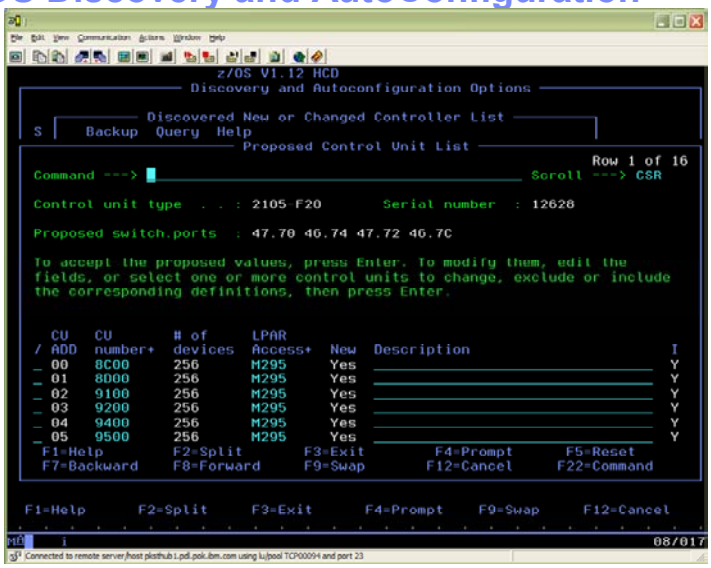

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration

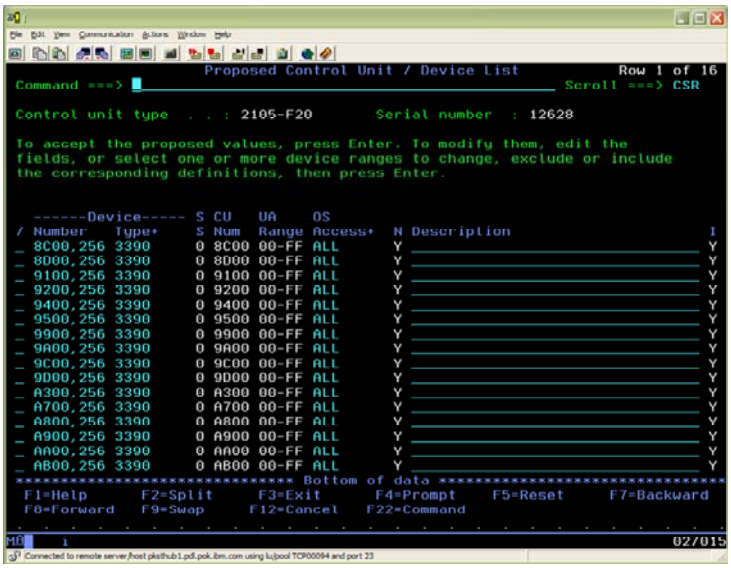

© 2010 IBM Corporation

z/OS Discovery and AutoConfiguration

- **Hardware Dependencies**
  - zDAC uses new capabilities in the z196 processor for fabric discovery
  - Explores FICON channels attached to switches
    - FICON Express8 or FICON Express4 (CHPID type: FC)
  - Dynamic I/O Enabled Partitions
  - Up to date controller microcode designed to improve discovery performance
    - Tested with IBM 2107 Controller level R12p.9b090910b - Bundle 64.30.87.0
  - Controller that registers with the fabric name server
    - All IBM FICON-attached tape and disk supported
  - Up to date switch microcode (see notes)
    - Tested with:
      - Brocade Firmware Version: V6.2.0e
      - McData Firmware Version : 09.09.00
- **Software Dependencies**
  - All target systems must be at z/OS R1.12 or higher
    - Mixed level SYSPLEXes are allowed, but only R1.12 can participate in discovery processing
  - Suggested use of Dynamic Channel path Management (DCM) for performance

© 2010 IBM Corporation

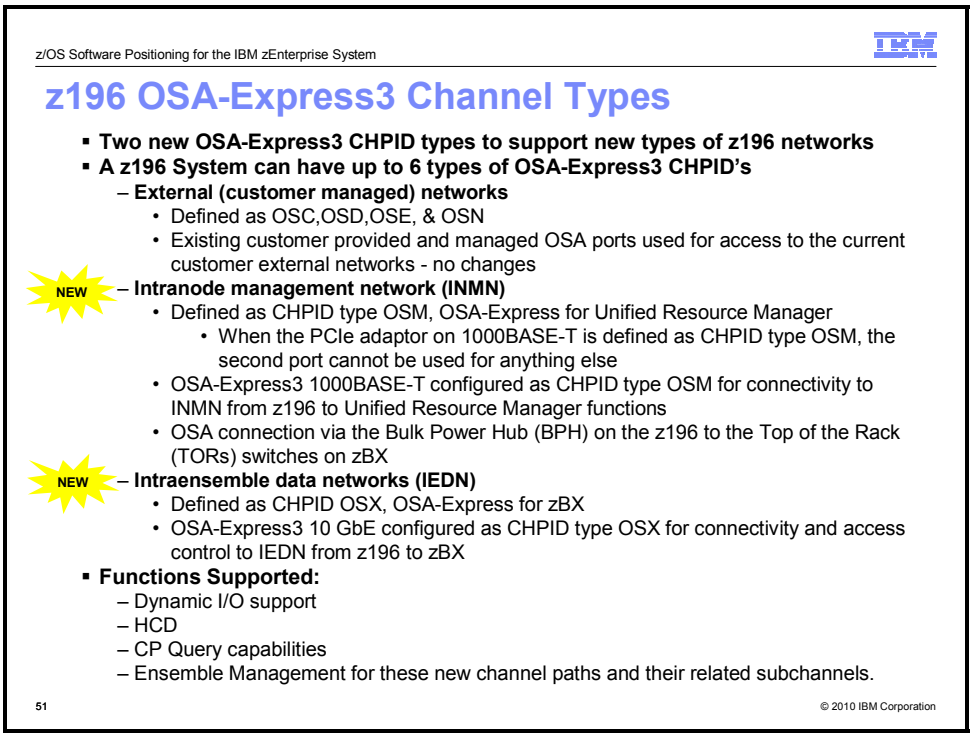

New OSA-Express3 Channel Types

The new zEnterprise System unifies the zEnterprise 196 server and the zEnterprise BladeCenter Extension, through the use of two new "internal" networks. Configured through OSA-Express3 adapters, the intranode management network (INMN) is used to provide system management functions and the intraensemble data network (IEDN) can be used to send data between the systems.

- The intranode management network (INMN) is a private and physically isolated OSA-Express3 1000BASE-T Ethernet internal management network operating at 1 Gbps, that connects all zEnterprise System resources (CPCs, Blade Centers, 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 your application data traffic. The INMN is exclusive to zEnterprise System, and is configured with a CHPID type of OSM.
- Intraensemble data network (IEDN) is a private OSA-Express3 10 Gigabit Ethernet 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 customer-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 zEnterprise System, and is configured with a CHPID type of OSX.

IBM zEnterprise Unified Resource Manager

IBM zEnterprise 196 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 z196 and its optionally attached IBM zEnterprise BladeCenter Extension (zBX) model 002. An ensemble can consist of a single z196 with no zBX attached, using an OSA loopback cable, or two to eight CPCs where at least one of the z196s 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**.
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, 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.

Ensemble membership requires that the z196 have two OSA-Express3 1000BASE-T features defined as CHPID type OSM for connectivity to the intranode management network (INMN). Communication across the INMN is exclusively for the purpose of enabling the Unified Resource Manager of the Hardware Management Console (HMC) to perform its various management disciplines (such as performance management, network virtualization management, energy management, etc.).

z196 OSA-Express3 CHPID Types

CHPID type	Purpose / Traffic	Operating Systems
OSD **All OSA features** **z196, z10, z9, zSeries**	Supports Queue Direct Input/Output (QDIO) architecture TCP/IP traffic when Layer 3 (uses IP address) Protocol-independent when Layer 2 (uses MAC address)	z/OS, z/VM z/VSE, z/TPF Linux on System z
OSE **1000BASE-T** **z196, z10, z9, zSeries**	Non-QDIO; for SNA/APPN/HPR traffic and TCP/IP "passthru" traffic	z/OS, z/VM z/VSE
OSC **1000BASE-T** **z196, z10, z9, z990, z890**	OSA-Integrated Console Controller (OSA-ICC) Supports TN3270E, non-SNA DFT to IPL CPCs & LPs	z/OS, z/VM z/VSE
OSM **1000BASE-T** **z196 exclusive**	OSA-Express for Unified Resource Manager Connectivity to intranode management network (INMN) from z196 to Unified Resource Manager functions	z/OS, z/VM Linux on System z*
OSN **GbE, 1000BASE-T** **z196, z10, z9** **exclusive**	OSA-Express for NCP Appears to OS as a device supporting CDLC protocol Enables Network Control Program (NCP) channel-related functions Provides LP-to-LP connectivity OS to IBM Communication Controller for Linux (CCL)	z/OS, z/VM z/VSE, z/TPF Linux on System z
OSX **10 GbE** **z196 exclusive**	OSA-Express for zBX Connectivity and access control to intraensemble data network (IEDN) from z196 to zBX	z/OS, z/VM Linux on System z*

52

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

© 2010 IBM Corporation

z196 Ensemble 'Network' Configuration

- **Each Ensemble requires**
  - Customer managed Management Network
    - A pair of (Primary & Alternate/Backup) HMCs for the Unified Resource Manager.
    - To control and manage the Ensemble.
    - Only one pair of Unified Resource Manager HMCs per Ensemble.
    - For high availability, one is a mirrored backup, and cannot be used as a Unified Resource Manager HMC, until fail-over
    - Both are attached to the Customer Managed Management Network.
  - Intranode management network
    - 2 ports from 2 different OSA-Express3 1000BASE-T Ethernet adapters, for redundancy.
    - To allow the Unified Resource Manager HMC, to talk to the System z hypervisors, within the managed Ensemble
- **Depending on the System z applications, each ensemble might require: (these are required if there is a zBX present)**
  - Intranode management network
    - Used for internal node hardware management. It is provided and managed by IBM
    - Connects the CPC's BHPs to the management port of the customer procured switch or the TOR switches in the zBX
  - Intraensemble data network
    - A pair of OSA-Express-3 10 GbE adapters, for redundancy.
    - To allow the System z applications to communicate between OS images to share data
    - To allow the System z application to communicate to the zBX
    - Ensemble zBX to zBX communications.
  - Customer managed data network
    - A pair of 10 Gb connections
    - Must provide the MAC address for each device
- **External customer networks (example OSC, OSD, OSE, and OSN connections) are outside the ensemble, and not included**

© 2010 IBM Corporation

Usage & Invocation (HCD)

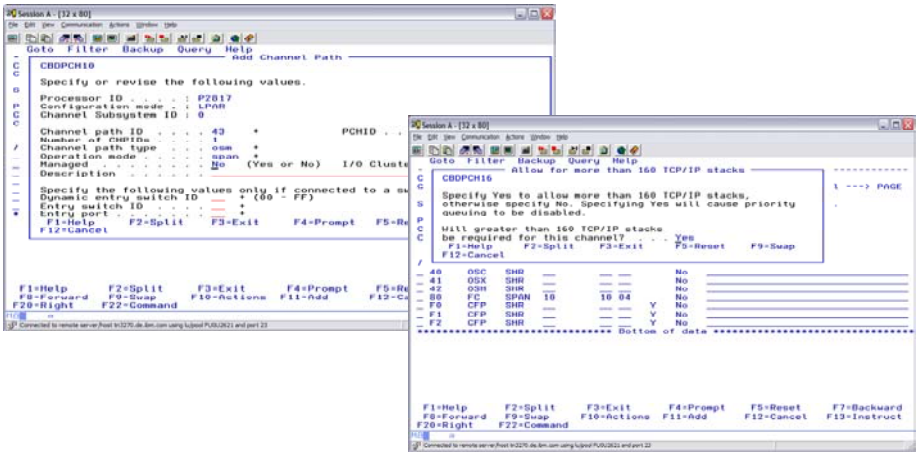

© 2010 IBM Corporation

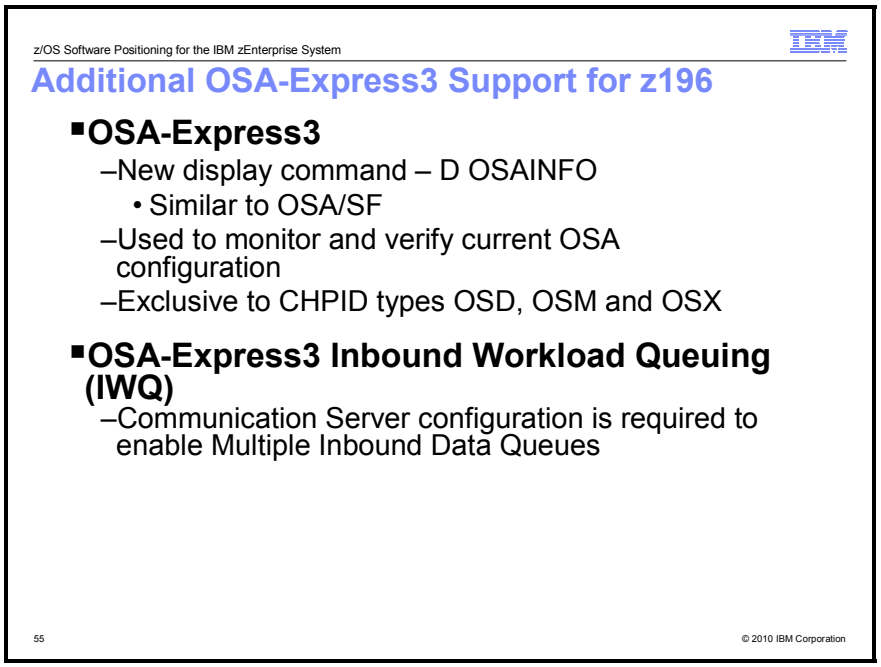

For Network Performance- Inbound Workload Queuing:
OSA-Express3 introduces inbound workload queuing (IWQ), which creates multiple input queues and allows OSA to differentiate workloads "off the wire" and then assign work to a specific input queue (per device) to z/OS. With each input queue representing 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.

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.

Customization / Activation zManager Performance Management

- Create ensemble and add CPC to ensemble
- Configure GPMP
  - The guest platform management provider (GPMP) is the interface between the *z196* intranode management network (INMN) and the z/OS Workload Manager
  - GPMP provides policy information to WLM about the platform wide performance goals of workloads in which the z/OS is participating
  - WLM supports the GPMP by:
    - configuring the GPMP with the WLM ISPF Application
    - managing the GPMP address space (start, stop, and restart)
    - displaying GPMP status information
    - collecting and aggregating performance measurements
- Start /Stop GPMP

© 2010 IBM Corporation

GPMP Configuration and Management

- On **policy activation**
  - WLM checks whether the service definition has valid GPMP settings
  - If activate=yes and system name not specified on excluded-list, GPMP is started automatically
- Also, you can use the **MODIFY WLM** command.The syntax of the command is:
  - **F WLM,GPMP,START**
    - Indicates that you want to start the GPMP
  - **F WLM,GPMP,STOP**
    - Indicates that WLM stops the currently active GPMP instance
  - **F WLM,GPMP,TRACE=NONE|LOW|MEDIUM|HIGH,DEST=FILE|MEMORY**
    - Enables you to change the GPMP internal tracing level “on the fly” and to change the destination of the trace (file or memory)
- Once you stopped the GPMP manually, the GPMP switches into “manual mode”. It is not automatically restarted even if a WLM policy with a valid GPMP configuration gets activated

© 2010 IBM Corporation

Concerning the **START** command:

- If the GPMP is already active, message IWM077E is issued.
- If ARM is disabled or the h/w does not support the GPMP, message IWM078E is issued.

Concerning the **STOP** command:

- WLM does not automatically restart the GPMP even if the WLM policy contains a valid GPMP configuration (i.e. activate=yes and system not in excluded-list). The GPMP is restarted only if the MODIFY WLM,GPMP,START command is issued from the MVS console. Even in the case of an activation of a new policy that contains a valid GPMP configuration, it is not restarted. The DISPLAY WLM,SYSTEMS operator command indicates that situation with a status of “STOPPED”.
- When the GPMP terminates unexpectedly or when too many errors occur, the GPMP state in the DISPLAY WLM,SYSTEMS command is set to “FAILED”.

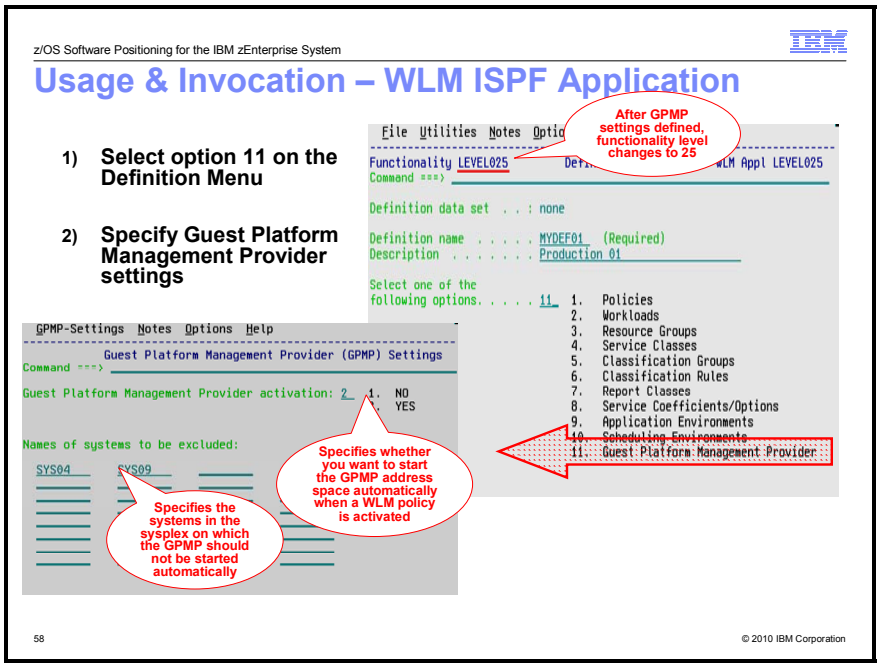

After selecting the new option 11 on the Definition Menu, you can specify whether you want to start the GPMP address space automatically when a WLM policy is activated. The valid options are YES or NO. In addition, you can specify the systems in the sysplex on which the GPMP should not be started automatically. These systems are excluded from the GPMP configuration.

Alternatively, if you have no service definition where the GPMP is activated, you can also start and stop the GPMP manually from the console.

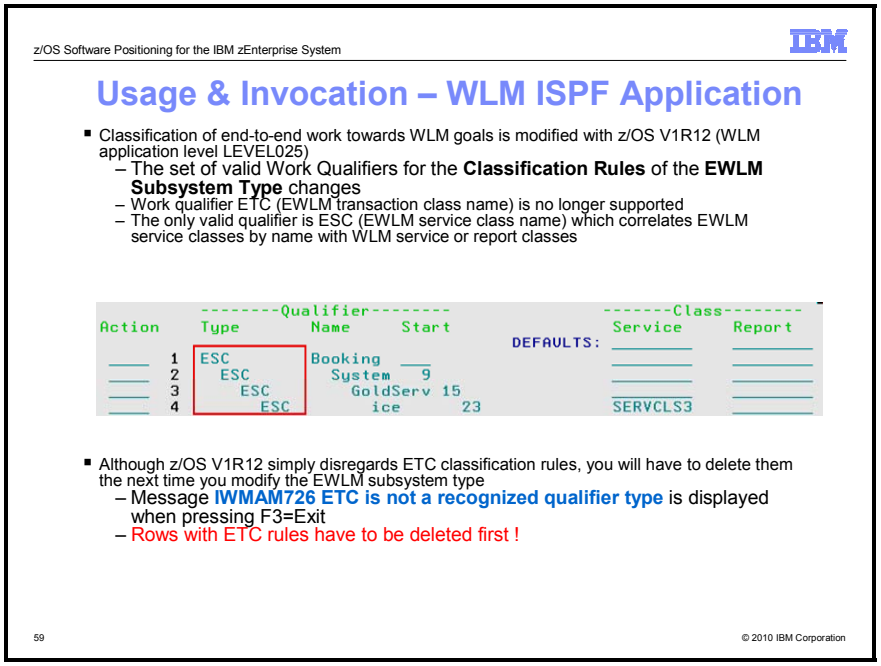

Beginning with z/OS V1R12, the workload management (WLM) service definition no longer supports the work qualifier EWLM transaction class name (ETC) for classification rules of the subsystem type EWLM. Although z/OS V1R12 disregards classification rules with the ETC work qualifier, you should consider removing them. If you do not remove the rules, you will have to delete them the next time you use the WLM ISPF application to modify the EWLM subsystem type.

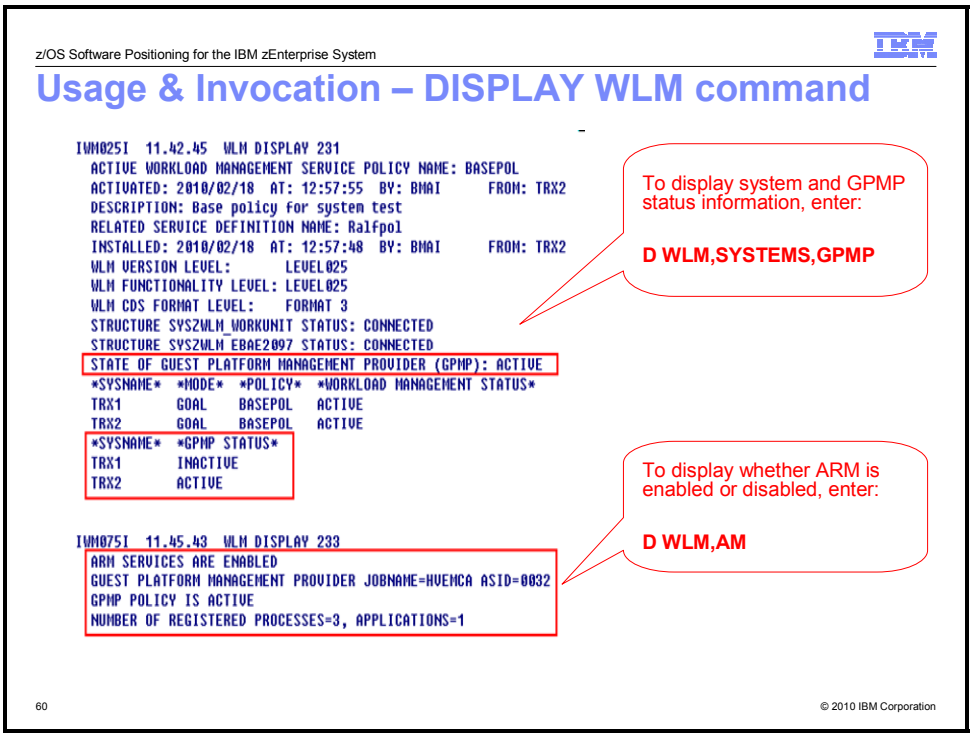

The **Display Workload Manager Information** command is extended to provide ARM and GPMP status information.

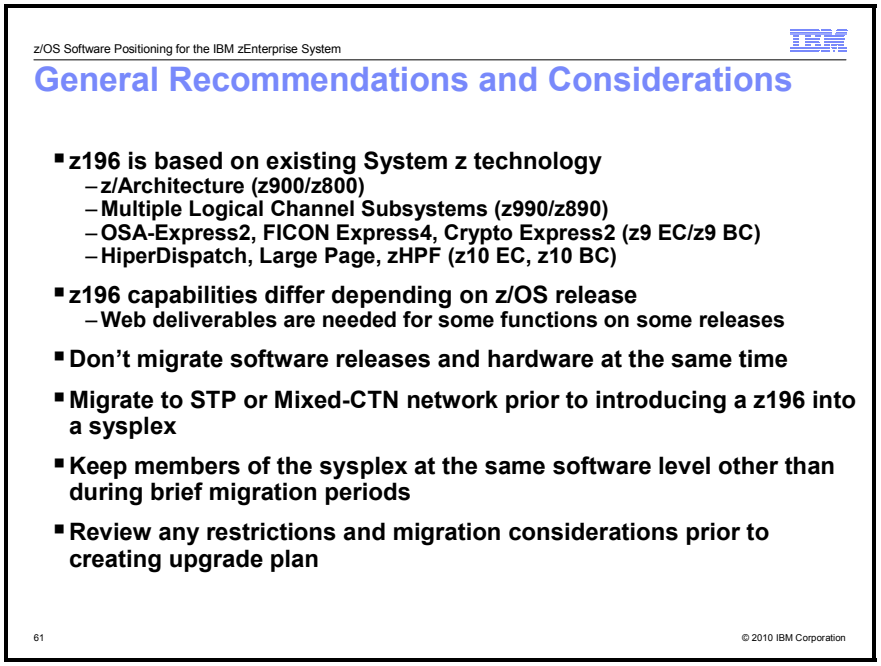

General Recommendations and Considerations

As you plan your migration to a z196 server, consider the following:

1. **Relatively few migration actions are new when coming from a z10 EC or a z10 BC server.** Migration to a z196 server has, as its base, a migration to the z10 EC or z10 BC servers. This means that if you are migrating to a z196 server from a z10 EC or z10 BC server, and have performed the migration actions associated with the z10 EC or z10 BC, you have fewer migration actions than if you were migrating from a server prior to the z10 EC or z10 BC and have not yet performed the migration actions associated with these servers. There are, in fact, very few new migration actions to perform on z/OS for a z196 server if you have already migrated to a z10 EC or z10 BC server. It is important to note that you can migrate directly to a z196 server without installing the intermediate (prior to z10 EC and z10 BC) servers, but you still need to ensure that any migration considerations are satisfied for the servers that you "skipped."
2. **Support is delivered by service (and FMID Web deliverables for ICSF).** The delta (from a z10 EC or z10 BC) support for a z196 server, excluding cryptographic support, is delivered by service (PTFs). The cryptographic support for the z196 server continues to be FMIDs, many of which are still available in Web deliverables. Different ICSF Web deliverables, providing different levels of support, are available for different releases of z/OS.
3. **Use the same software level throughout a sysplex.** Having members of a sysplex at the same software level (other than during brief migration periods) is good software management policy.
4. **Migrate hardware and software at different times.** To minimize the amount of change (and therefore risk) that you experience at one time, do not migrate your software release level at the same time that you migrate your hardware.
5. **Review the sysplex configuration in which the z196 server will participate.** In particular, if you have any existing z900, z800, z990, or z890 z/OS images or coupling facility images in the sysplex, move those z/OS images or coupling facilities to later servers, such as a System z10 (z10 EC or z10 BC), System z9 (z9 EC or z9 BC), or later. This action is necessitated by the restriction that a z196 server cannot participate with prior-level servers in a sysplex.
6. **Implement STP (or a Mixed-CTN) timing network.** This action is necessitated because Sysplex Timers (9037-002) are not supported on a z196 server.
7. **Migrate from ICB-4 to InfiniBand coupling links.** This action is necessitated because ICB-4 links are not supported on z196 servers. If desired, you can take this action after you order a z196, as you upgrade to the new server.
8. **Migrate from unsupported hardware features to newer technology.** This action is necessitated because FICON Express, FICON Express2, Crypto Express2, and OSA-Express2 10 GbE LR are not supported on z196 servers.

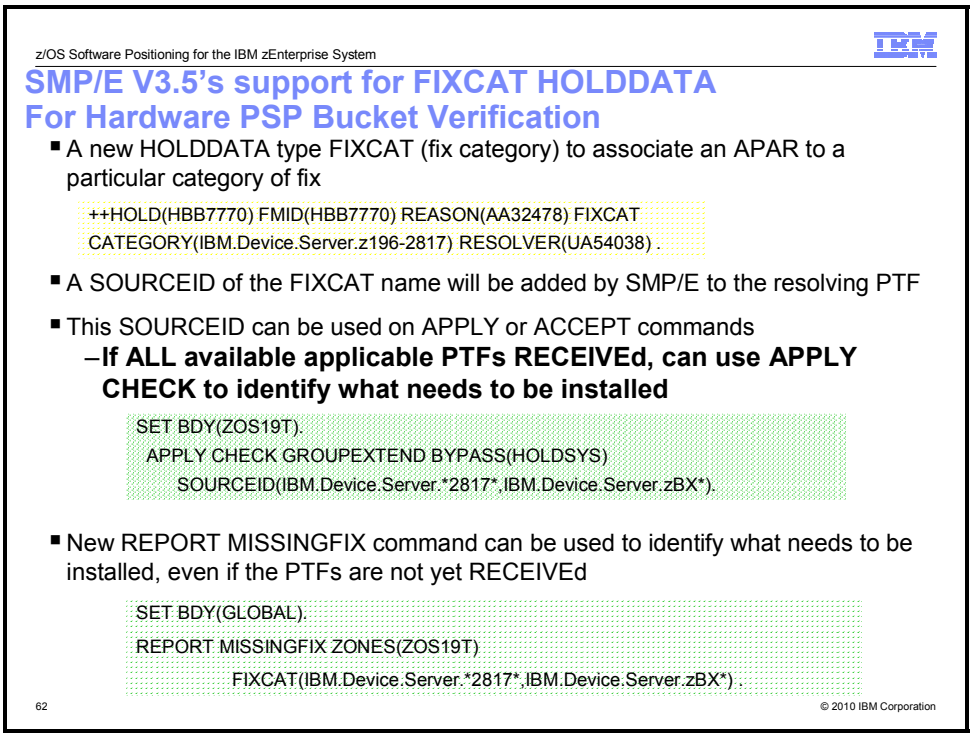

PSP Buckets

For a z196 CPC, PTFs are identified in the 2817DEVICE PSP bucket (Subset 2817/ZOS). For an IBM zEnterprise BladeCenter Extension (zBX) attached to your z196 CPC, either to support IBM Smart Analytics Optimizer for DB2 for z/OS, V1.1 (5697-AQT) or to support IBM POWER7 Blades, the PTFs are identified in the 2458DEVICE PSP bucket (Subset 2458/ZOS). In each PSP bucket, the content is dependent on the z/OS release you will run on the z196 server. If you reviewed the PSP buckets some time ago, review them again to ensure that any newly identified z/OS service has been installed. To assist you in determining if you have the recommended service (identified in these PSP buckets) installed on your system, you can use the SMP/E REPORT MISSINGFIX command in conjunction with the FIXCAT type of HOLDDATA, as follows:

a. Acquire and RECEIVE the latest HOLDDATA onto your z/OS system(s). Use your normal service acquisition portals or download the two (2) year HOLDDATA directly from http://service.software.ibm.com/holdata/390holddata.html. Ensure you select **Full** from the Download NOW column (last 730 days) to receive the FIXCAT HOLDDATA, as the other files do not contain FIXCAT HOLDDATA.
b. Run the SMP/E REPORT MISSINGFIX command on your z/OS systems and specify one or more of the following Fix Categories (FIXCAT):
   - IBM.Device.Server.z196-2817
   - IBM.Device.Server.z196-2817.ParallelSysplexInfiniBandCoupling
   - IBM.Device.Server.z196-2817.ServerTimeProtocol
   - IBM.Device.Server.z196-2817.zHighPerformanceFICON
   - IBM.Device.Server.zBX-2458
   - IBM.Device.Server.zBX-2458.ISAOPT

   The report will identify any missing coexistence and fallback PTFs for that system. For complete information about the REPORT MISSINGFIX command, see *SMP/E Commands*.
c. Periodically, you might want to acquire the latest HOLDDATA and rerun the REPORT MISSINGFIX command to find out if there are any new PTFs recommended for the z196 server.

If you are skipping a server generations, review the PSP buckets for the servers you skipped, as follows:

- For System z10, see Upgrade 2097DEVICE, Subset 2097/ZOS.
- For System z9, see Upgrade 2094DEVICE, Subset 2094/ZOS/1.
- For System z990 or z890, see Upgrade 2084DEVICE, Subset 2084/ZOS/1.

SMP/E 3.5 (also part of z/OS V1.10 – which became availability September 26, 2008) provided a new function that can greatly simplify the identification, verification, acquisition, and installation tasks for service documented in PSP buckets by processing a new kind of HOLDDATA that associates PTFs to fix categories. Each ++HOLD statement identifies an APAR and its resolving PTF needed to provide support for the fix category. The new FIXCAT HOLDDATA is now being delivered with existing ERROR HOLDDATA. Fix categories have been defined for hardware devices, new software FMIDs, and selected hardware and software functions.

When SMP/E V3.5 processes FIXCAT ++HOLD statements (during RECEIVE processing) the resolving PTF identified on the ++HOLD statement is assigned an SMP/E SOURCEID that matches the fix category on that statement. This enables you to specify fix categories as part of the SOURCEID operand on APPLY, ACCEPT, and LIST commands.

Many customers fall into one of two camps when it comes to acquiring and installing service:
1. You choose to RECEIVE all available applicable PTFs and filter which PTFs get installed using parameters on the APPLY command; or
2. You only acquire PTFs you need when you want to install them

No matter which approach you use, SMP/E V3.5 can simplify your task.

In the first scenario, you have all available PTFs that are applicable to your environment already RECEIVEd. To install all the PTFs needed for minimal support of a z196 server, you could use something like the following command. You don't have to access the PSP bucket, nor download any extra files.

```
APPLY CHECK GROUPEXTEND BYPASS(HOLDSYS)
SOURCEID(IBM.Device.Server.z196-2817, IBM.Device.Server.zBX-2458).
```

If you wanted to include all support for the z196 optional functions, then the SOURCEID operand could either list all fix categories that you are interested in, or use wildcarding to have SMP/E determine all applicable fix categories. This latter approach would automatically express interest in future z196 functions (should new fix categories be created). Specifying IBM.Device.Server.z196-2817*, or IBM.Device.Server.z196*, or just *2817* all result in similar results: an interest in the minimal support for z196 as well as exploitation of all optional functions. In these examples the asterisk ('*') represents 0-n characters, while the percent sign ('%') could be used to represent a single character. Based on the structured names of the fix categories for hardware support, you can use SOURCEIDs with wildcarding to ensure that PTFs for multiple generations of servers are installed.

In the second scenario, you don't know if you have RECEIVEd all the PTFs that you need since you only obtain PTFs right before you want to install them. You need to first obtain the latest HOLDDATA using your favorite method (mine is RECEIVE ORDER(CONTENT(HOLDDATA))...). Then you can use the new REPORT MISSINGFIX command to identify PTFs missing from your system. For example:

```
SET BDY(GLOBAL).
REPORT MISSINGFIX ZONES(ZOSR12T)
     FIXCAT(IBM.Device.Server.z196-2817, IBM.Device.Server.zBX-2458) .
```

The REPORT MISSINGFIX command checks your GLOBAL zone for FIXCAT HOLDDATA that matches the FIXCAT values specified on the command. The command then compares the APARs identified in that FIXCAT HOLDDATA with the PTFs installed in the specified zones, and produces a report to identify any APARs that are not resolved. In other words, it reports which PTFs (fixes) are missing for the specified fix categories. Furthermore, the command produces a customized job that can be used to obtain any PTFs not already RECEIVEd via the RECEIVE ORDER command, and install any missing service via the APPLY CHECK command. Please note that the FIXCAT operand on the REPORT MISSINGFIX command can list multiple fix categories, as well as using the same wildcarding techniques described above for the SOURCEID operand.

Because both of these techniques are simple and integrated into basic SMP/E commands, you can use them periodically to ensure that you have the latest PTFs specified in the hardware PSP bucket installed (since PSP buckets can be updated daily). You can even perform these same checks as you upgrade z/OS releases to ensure that you have the appropriate PTFs for your existing hardware installed for your new z/OS release.

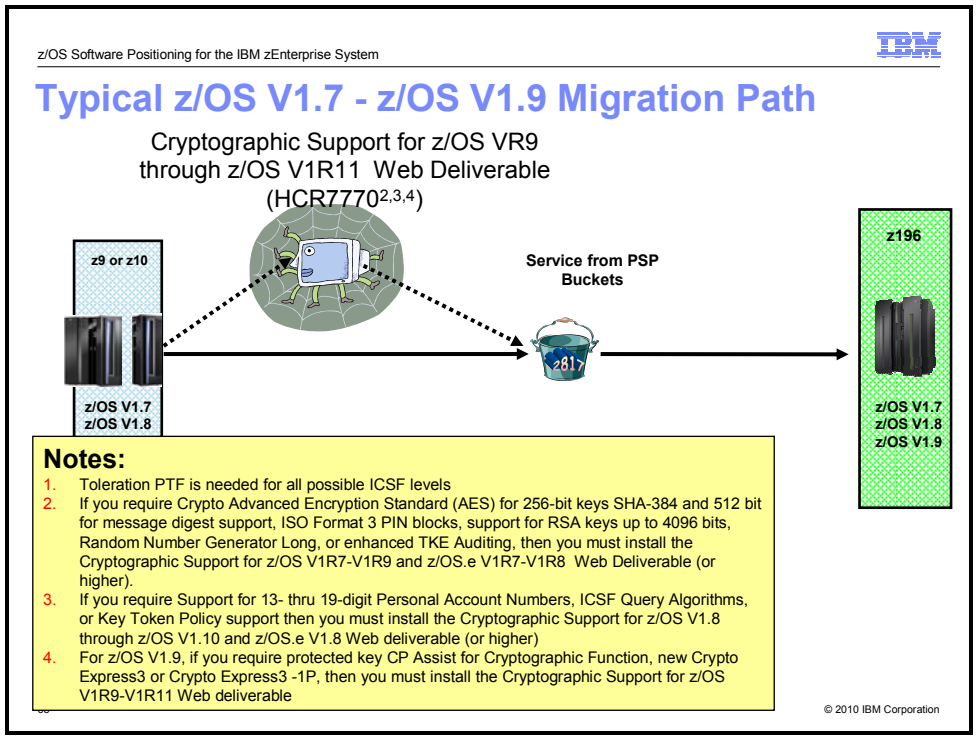

Migration steps if you are running z/OS V1.7, z/OS V1.8, or z/OS V1.9

The steps for migrating from z/OS V1.7, z/OS V1.8, or z/OS V1.9 on your current server to z/OS V1.7, z/OS V1.8, z/OS V1.9 on a z196 server, are:

1. Ensure that you are licensed for the IBM Lifecycle Extension for z/OS V1.7 (5637-A01), the IBM Lifecycle Extension for z/OS V1.8 (5638-A01), or the IBM Lifecycle Extension for z/OS V1.9 (required as of October 2010). With the Lifecycle Extension, z/OS V1.7, V1.8, and V1.9 support the z196 server.
2. Install the necessary z/OS service, as indicated in PSP buckets. For a z196 CPC, PTFs are identified in the 2817DEVICE PSP bucket (Subset 2817/ZOS). For an IBM zEnterprise BladeCenter Extension (zBX) attached to your z196 CPC, either to support IBM Smart Analytics Optimizer for DB2 for z/OS, V1.1 (5697-AQT) or to support IBM POWER7 Blades, the PTFs are identified in the 2458DEVICE PSP bucket (Subset 2458/ZOS). In each PSP bucket, the content is dependent on the z/OS release you will run on the z196 server.
3. Ensure that you have the *IBM zIIP Support for z/OS and z/OS.e V1R6/V1R7* Web Deliverable installed if you are running z/OS V1.7. You do not need to have a zIIP processor, only the FMID installed. This will enable you to use HiperDispatch. If you don't want to use HiperDispatch, you don't need to have the Web deliverable installed.
4. Determine the level of cryptographic function that you want to use. Cryptographic toleration support is provided by PTFs installed on one of the Web deliverables listed below. The support varies depending on the Web deliverable that is installed. Note that you must also install the PTFs identified in the program directory that comes with the Web deliverables and the PTFs that are recommended in the appropriate ICSF PSP bucket.
   - If you are running z/OS V1.7, minimally you will have to install toleration PTFs on top of the *Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/R7* Web deliverable (FMID HCR7731).
   - Install the *Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8* Web deliverable (HCR7750) or higher if you are running z/OS V1R7, z/OS V1R8, or z/OS V1R9 and you require Crypto Advanced Encryption Standard (AES) for 256-bit keys SHA-384 and 512-bit for message digest support, ISO Format 3 PIN blocks, support for RSA keys up to 4096 bits, Random Number Generator Long, or enhanced TKE Auditing.
   - Install the *Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8* Web deliverable (HCR7751) or higher if you are running z/OS V1R7, z/OS V1R8, and z/OS V1R9 and you require

Support for 13- thru 19-digit Personal Account Numbers, ICSF Query Algorithms, or Key Token Policy support.

- Install the *Cryptographic Support for z/OS V1R9-V1R11 Web* deliverable (HCR7770) or higher if you are running z/OS V1R9 and you require protected key CP Assist for Cryptographic Function, new Crypto Express3 or Crypto Express3 -1P.

Please note

1. APARs/PTFs are required to support the Cryptographic Express3 Coprocessor (CEx3C). These were also required on a z10 server that had a Cryptographic Express3 Coprocessor. Alternatively for z/OS V1.9, you can install the *Cryptographic Support for z/OS V1R9-V1R11* Web deliverable (FMID HCR7770).
2. If you are sharing keys with a z/OS V1.10 or higher system that is using the Cryptographic Support for z/OS V1R10-V1R12 Web Deliverable (FMID HCR7780), then the z/OS V1.7-z/OS V1.9 systems will need coexistence PTFs. To assist in identifying the coexistence service, you could use the Fix Category: IBM.Coexistence.ICSF.z/OS_V1R10-V1R12-HCR7780

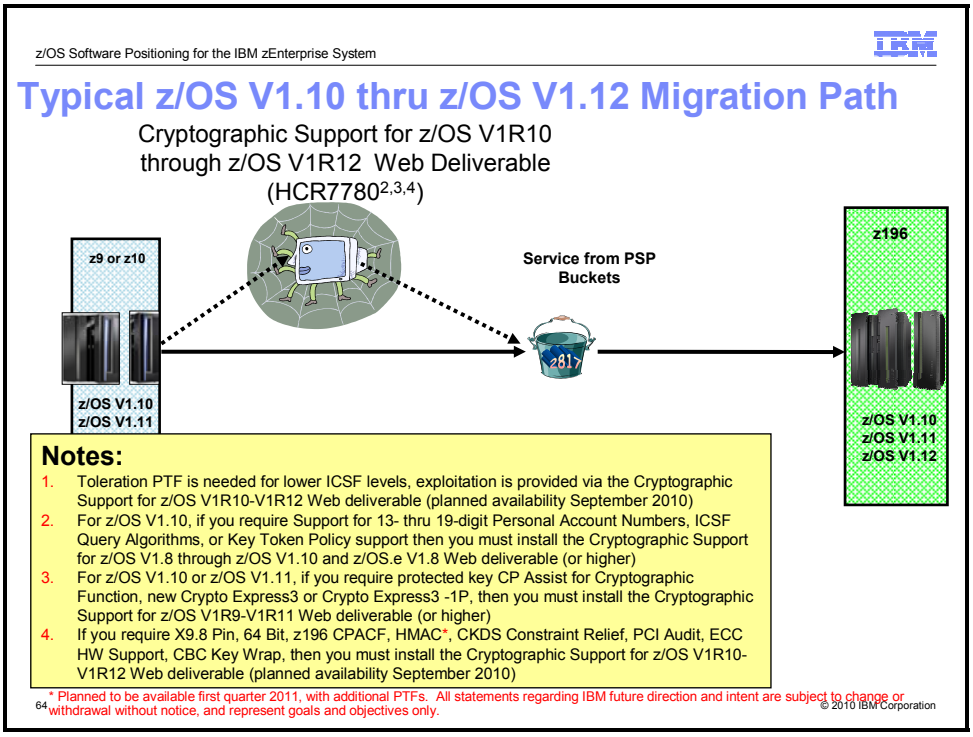

Migration steps if you are running z/OS V1.10 or z/OS V1.11 or z/OS V1.12

The steps for migrating from z/OS V1.10, z/OS V1.11, or z/OS V1.12 on your current server to z/OS V1.10, z/OS V1.11, or z/OS V1.12 on a z196 server are:

1. Install the necessary z/OS service, as indicated in PSP buckets. For a z196 CPC, PTFs are identified in the 2817DEVICE PSP bucket (Subset 2817/ZOS). For an IBM zEnterprise BladeCenter Extension (zBX) attached to your z196 CPC, either to support IBM Smart Analytics Optimizer for DB2 for z/OS, V1.1 (5697-AQT) or to support IBM POWER7 Blades, the PTFs are identified in the 2458DEVICE PSP bucket (Subset 2458/ZOS). In each PSP bucket, the content is dependent on the z/OS release you will run on the z196 server.

2. Determine the level of cryptographic function that you want to use. Cryptographic toleration support is provided by PTFs installed on one of the Web deliverables listed below. The support varies depending on the Web deliverable that is installed. Note that you must also install the PTFs identified in the program directory that comes with the Web deliverables and the PTFs that are recommended in the appropriate ICSF PSP bucket.
   - Install the *Cryptographic Support for z/OS V1R8-V1R10* Web deliverable (FMID HCR7751) or higher if you are running z/OS V1.10 and require support for 13- through 19-digit Personal Account Numbers, ICSF Query Algorithms, or Key Token Policy support.
   - Install the *Cryptographic Support for z/OS V1R9-V1R11* Web deliverable (FMID HCR7770) or higher if you are running z/OS V1.10 or z/OS V1.11 and you require protected key CP Assist for Cryptographic Function, new Crypto Express3 or Crypto Express3 -1P.
   - Install the *Cryptographic Support for z/OS V1R10-V1R12* Web deliverable (FMID HCR7780) if you are running z/OS V1.10, z/OS V1.11, or z/OS V1.12 and you require X9.8 Pin, 64 Bit, z196 CPACF, HMAC, CKDS Constraint Relief, PCI Audit, ECC HW Support, CBC Key Wrap. The planned availability of this Web deliverable is September 2010. The planned availability of HMAC support is first quarter 2011, with additional PTFs.

   Please note
   1. For z/OS V1.10 and z/OS V1.11 APARs/PTFs are required to support the Cryptographic Express3 Coprocessor (CEx3C). These were also required on a z10 server that had a Cryptographic Express3 Coprocessor. Alternatively, you can install either the *Cryptographic Support for z/OS V1R9-V1R11* Web deliverable (FMID HCR7770) or the *Cryptographic Support for z/OS V1R10-V1R12* Web deliverable (FMID HCR7780).
   2. If you are sharing keys with a z/OS V1.10 or higher system that is using the Cryptographic Support for z/OS V1R10-V1R12 Web Deliverable (FMID HCR7780), then the z/OS V1.10-z/OS V1.12 systems will need coexistence PTFs. To assist in identifying the coexistence service, you could use the Fix Category: IBM.Coexistence.ICSF.z/OS_V1R10-V1R12-HCR7780

3. You can define an Ensemble and exploit the IBM zEnterprise Unified Resource Manager. See *System z Ensemble Planning and Configuration Guide* (GC27-2608) for a detailed description of the steps required to define an Ensemble.

© 2010 IBM Corporation

Summary

© 2010 IBM Corporation

Supported IBM zEnterprise System Migration Scenarios

1. Existing System z server to new z196 server
   - Toleration versus exploitation differs by release
   - ICSF web deliverable needed for function required
2. Existing z9 or z10 sysplex replaced by z196 (or z196 added to existing sysplex)
   - Coupling Link technology (no ICB-4 to z196)
   - Only STP or Mixed CTN supported (no ETR connections to a z196)
3. All existing z9 or z10 z/OS images replaced by z196 images
   - If you want an ensemble
     - z/OS releases must be z/OS V1.10 or higher
     - Configure OSM OSA CHPIDs for Intranode Management Network (INMN)
4. Exploit a z196 with zBX using IBM Blades
   - z/OS releases must be z/OS V1.10 or higher
   - Configure OSM OSA CHPIDs for Intranode Management Network (INMN)
   - Configure OSX OSA CHPIDs for Intraensemble Data Networks (IEDN)
   - User responsible for AIX OS and application provisioning
5. Exploit z196 with zBX using IBM Smart Analytics Optimizer
   - z/OS releases must be z/OS V1.10 or higher; DB2 releases V9.1 or higher
   - Configure OSM OSA CHPIDs for Intranode Management Network (INMN)
   - Configure OSX OSA CHPIDs for Intraensemble Data Networks (IEDN)

© 2010 IBM Corporation

Summary: z/OS Software Support for IBM zEnterprise 196

- z/OS V1.7, z/OS V1.8 and z/OS V1.9
  - New MSU levels
  - New z/Architecture Instructions (and new OPCODE support)
  - IFAURP Reporting
  - Toleration for >64 CPs on the server
  - Crypto toleration
- z/OS V1.10 and z/OS V1.11
  - zBX Support for optimizers and IBM POWER7 blades
  - IBM zEnterprise Unified Resource Manager
    - Network Management
      - Intranode Management Network (INMN)
      - Intraensemble data network (IEDN)
    - Performance Management
  - RSM exploitation of new OPCODEs
  - HiperDispatch cache and affinity node changes
  - HiperDispatch performance and serviceability enhancements
  - Static Power Save Mode
  - Language Environment High Register Resolution
  - CPU Measurement Facility (Hardware Instrumentation Services)
  - CF Level 17
    - CFCC Diagnostics
    - Scalability improvements with up to 80 coupling links
    - Connectivity improvements with up to 2047 structures
    - Connectivity improvements with 128 CHPIDs per server
  - Removal of the 64k byte data transfer limit for zHPF multitrack operations
  - Three subchannel sets
  - Crypto Exploitation (ANSI X9.8 Pin security, enhanced Common Cryptographic Architecture (CCA), 64 Bit, CP Assist for Cryptographic Function (CPACF) enhancements, Secure Keyed-Hash Message Authentication Code (HMAC)*, CKDS Constraint Relief, PCI Audit, Elliptical Curve Cryptography (ECC) Digital Signature Algorithm, and CBC Key Wrap)
  - RMF Postprocessor Crypto Activity report support for 4096-bit
- z/OS V1.12
  - XL C/C++ ARCH(9) and Tune(9) Options
  - zDAC support
  - New display command – D OSAINFO
  - OSA-Express3 Inbound Workload Queuing (IWQ)

* Planned to be available first quarter 2011, with additional PTFs. All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

© 2010 IBM Corporation

z196 PSP Bucket – 2817DEVICE 2817/ZOS

z/OS Software Release	Base Support																	Explicit z/OS Support											
	Base GA1 Support	OSA-Express3 (GbE LX andx SX, 1000BASE-T, 10 GbE LR and SR)	New z/Architecture Instructions[3]	IB Coupling Links	Up to 128 Coupling Link CHPIDs	FICON Express8 (CHPID FC)	IFAURP Reporting	>64 CPs per Server	Crypto Toleration[4]	HiperDispatch	OSA-Express3 CHPID OSD maximum ports	RMF Postprocessor Crypto Activity report - 4096-bit[6]	CPU Measurement Facility (HIS)	> 64 CPs per LPAR	HiperDispatch Cache and Affinity Node Changes	HiperDispatch Serviceability	LE High Register Resolution	Power Save Mode	CF Level 17	Extensions to zHPF Multitrack Operations	Three Subchannel Sets	Crypto Exploitation	Up to 32 HiperSockets	IBM zEnterprise Blade center Extension (zBX) support	Network Management	Performance Management	zDAC Support	New OSA Display Command	OSA-Express3 Inbound Workload Queuing (IWQ)	X/L C/C++ ARCH(9) Tune(9)
z/OS V1.7[1]	P	P	P	P	P	P	P	P	WP	N	N	P[2]	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS V1.7[1] w/ZIIP	P	P	P	P	P	P	P	P	WP	WP	N	P[2]	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS V1.8[1]	P	P	P	P	P	P	P	P	P	BP	P	P	P	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS V1.9[2]	P	P	P	P	P	P	P	P	P	BP	P	P	P	N	N	N	N	P[5]	N	N	N	N	N	N	N	N	N	N	N
z/OS V1.10	P	P	P	P	P	P	P	P	P	B	B	P	P	P	P	P[1]	P	P	P	P	P	WP	P	P	P	P	N	N	N	N
z/OS V 1.11	P	P	P	P	P	P	P	B	P	B	B	P	P	B	P	P[1]	P	P	P	P	P	WP	P	P	P	P	N	N	N	N
z/OS V 1.12	BP	BP	BP	B	BP	B	B	B	P	B	B	B	B	B	B	B	B	B	BP	B	B	WP	B	B	B	B	B	B	B	B

[1] – The Lifecycle Extension for z/OS V1.7 or z/OS V1.8 is required for support
[2] – The IBM Lifecycle Extension for z/OS V1.9 will be required after 9/30/2010
[3] – Support differs by release, z/OS V1.12 includes XL C/C++ support for ARCH(9) and TUNE(9) Options
[4] – A Crypto Web Deliverable is NOT required, but toleration PTF is needed even if a web deliverable is installed
[5] – Toleration PTFs needed to detect Power Saving Mode
[6] – PTFs recognize crypto control block, AND provide updated reports
B – FMIDin Base product
W – FMIDs shipped in a Web Deliverable

P – PTFs are required, P[1] – Support differs by release
P[2]- PTF required for toleration, N – Not Supported

© 2010 IBM Corporation

Backup Material

© 2010 IBM Corporation

Additional Information

- z/OS Home Page
  http://www.ibm.com/servers/eserver/zseries/zos/
- zFavorites for System z
  http://www.ibm.com/servers/eserver/zseries/zos/zfavorites/
- z/OS Internet Library
  http://www.ibm.com/servers/eserver/zseries/zos/bkserv/
- IBM System z
  http://www.ibm.com/systems/z/
- IBM Resource Link
  https://app-06.www.ibm.com/servers/resourcelink/hom03010.nsf
- IBM Redbooks – How-To Books (also Redpieces)
  http://www.redbooks.ibm.com/
- Preventive Service Planning buckets
  http://www14.software.ibm.com/webapp/set2/psp/srchBroker
- Enhanced HOLDDATA
  http://service.boulder.ibm.com/390holddata.html

© 2010 IBM Corporation

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BladeCenter*	HiperSockets	Language Environment*	Redbooks*	SystemPac*	System z10	z10 BC
DB2*	IBM*	Parallel Sysplex*	RMF	System Storage*	System z10 Business Class	z10 EC
ESCON*	IBM (logo)	POWER7	ServerPac*	System z	z/Architecture*	z/OS*
FICON Express*	Infiniband*	RACF*	Sysplex Timer*	System z9	z9*	zEnterprise
					z10	zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Firefox is a trademark of Mozilla Foundation

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Internet Explorer is a trademark of Microsoft Corp

InfiniBand is a trademark and service mark of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

See url http://www.ibm.com/legal/copytrade.shtml for a list of IBM trademarks.

© 2010 IBM Corporation

z/OS Support For z10 Servers

© 2010 IBM Corporation

z/OS Support Summary for z10 Support

- **z/OS V1.7 (with the IBM Lifecycle Extension for z/OS V1.7)**
  - Parallel Sysplex InfiniBand (PSIFB) Coupling Links
  - OSA-Express3 10 Gigabit Ethernet
  - 65535 MP Factors
  - Crypto Toleration
    - CP Assist for Cryptographic Functions and Configurable Crypto Express2, Remote loading of ATM and Point of Sale (POS) keys,…
  - Crypto Exploitation
    - Crypto Advanced Encryption Standard (AES) for 256-bit keys SHA-384 and 512 bit for message digest support, RSA keys up to 4096 bits in Length, Dynamically Add Crypto to Logical Partition, Support for 13-19 digit PANs data, Secure Key AES Support, …
  - HiperDispatch (Requires zIIP Web Deliverable to be installed (but not a zIIP processor)
  - Decimal Floating Point (High Level Assembler support only)
  - High Performance FICON
  - CF Level 16
- **z/OS V1.8 and z/OS.e V1.8 (with the IBM Lifecycle Extension for z/OS V1.8)**
  - Large Memory > 128GB (up to 1 TB on z10 EC)
  - Additional support for Decimal Floating Point (HLASM, LE, DBX, and CDA RTLE)
    - XL C/C++ support if you download and install the z/OS V1.9 XL C/C++ Compiler
  - OSA-Express3 Double Port Density
  - CPU Measurement Facility
  - Worldwide port name (WWPN) tool for advance SAN preplanning
- **z/OS V1.9**
  - Additional support for Decimal Floating Point (XL C/C++ Compiler)
  - z/OS XL C/C++ Exploitation of z10 EC Machine Instructions
  - Greater than 54 CPs for a single LPAR
  - Large Page Support
  - Capacity Provisioning
  - RMF FICON Enhancements
  - HiperSockets Multiple Write Facility
  - OSA-Express3 Double Port Density
  - STP Enhancements
- **z/OS V1.10**
  - Dynamic add logical CPs
- **z/OS V1.11**
  - New XL C/C++ Compiler Options
  - Large Page Support in Language Environment
  - BCPii Activation Profile Support

© 2010 IBM Corporation

z10 EC PSP Bucket – 2097DEVICE 2097/ZOS
z10 BC PSP Bucket – 2098DEVICE 2098/ZOS

	Included in Base z/OS Support											Explicit z/OS Support																		
Release	z10 EC GA1 Support	OSA-Express3 10 GbE	OSA-Express3 4-port OSA-ICC	65535 MP Factors	z10 mnemonic support	Crypto Toleration[1]	HiperDispatch	Large Memory > 128GB (1 TB)[2]	C/C++ ARCH(8) and TUNE(8)	Greater than 54 CPs (64)[3]	Dynam. Add Logical CPs	PSIFB Coupling Links	Usage Report Program	Decimal Floating Point[4]	High Performance FICON	CF Level 16	Crypto Exploitation[5]	Layer 3 VMAC Support	Service Aids Support for Large Dumps	CPU Measurement Facility	OSA-Express3 Double Port Density	WWPN prediction tool	Large Page Support	Capacity Provisioning	HiperSockets Multiple Write Facility	STP Enhancements	Protected Key CPACF	OSA Optimized Latency Mode (OLM)	OSA QDIO data connection isolation	New XL C/C++ Compiler Options	BCPii Activation Profile
z/OS 1.7[6]	P	B	B	P	P	WP	N	N	N	N	N	P	P	P	P	P	W	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS 1.7[6] w/ZIIP	P	B	B	P	P	WP	WP	N	N	N	N	P	P	P	P	P	W	N	N	N	N	N	N	N	N	N	N	N	N	N
z/OS 1.8[6]	P	B	B	P	P	BP	P	BP	W	N	N	P	P	WP	P	P	W	P	P	P	P	P	N	N	N	N	N	N	N	N
z/OS 1.9	P	B	B	B	P	BP	BP	B	BP	B	N	BP	P	BP	P	P	W	B	P	P	P	P	BP	BP	P	P	WP	N	N	N	N
z/OS 1.10	BP	B	B	B	B	BP	BP	B	B	B	B	BP	P	B	P	P	W	B	B	P	B	P	BP	BP	B	P	WP	N	N	N	N
z/OS 1.11	BP	B	B	B	B	BP	B	B	B	B	B	B	B	B	B	B	W	B	B	B	B	B	B	B	B	B	B	WP	P	P	P

1 – New cryptographic toleration required for GA3 (Y4)
2 - Up to 1 Terabyte on a z10 EC, and up to 248 GB on a z10 BC
3 – Only on a z10 EC
4 – Decimal Floating Point Support differs by release
5 – Crypto Exploitation differs by release and web deliverable installed
6 - The IBM Lifecycle Extension for z/OS is required after Sept. 2009

© 2010 IBM Corporation

ICSF

© 2010 IBM Corporation

Operating System	Level Shipped in Base Product	Level Required for Secure or Clear Key or Enhanced Secure Key Support and 64 bit adressing caller support	Level Required for AES, PRNG, and SHA-256 and CEX2n exploitation	Level required for Remote Key Loading for ATMs and POSs	Level Required for 4096-RSA keys, ISO-3 PIN Support	Level Required for CPACF support for SHA-384 and SHA-512 Support	Level Required for 13- thru 19-digit Personal Account Numbers, ICSF Query Algorithms, Secure Key AES and Key Token Policy support	•Crypto Express3, Crypto Express3 1P, and Protected Key CPACF	X9.8 Pin, 64 Bit, z196 CPACF, HMAC*, CKDS Constraint Relief, AP Interrupt, PCI Audit, ECC HW Support, CBC Key Wrap
z/OS 1.7[1] or z/OS.e 1.7	HCR7720	Integrated in z/OS base level	HCR7730	HCR7731	HCR7750	HCR7750	HCR7751	N/A	N/A
z/OS 1.8[1] or z/OS.e 1.8	HCR7731	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	HCR7750	HCR7750	HCR7751	N/A	N/A
z/OS 1.9[1]	HCR7740	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	HCR7750	HCR7750	HCR7751	HCR7770	N/A
z/OS 1.10	HCR7750	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	HCR7751	HCR7770	HCR7780
z/OS 1.11	HCR7751	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	HCR7770	HCR7780
z/OS V1.12	HCR7770	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	Integrated in z/OS base level	HCR7780

* Planned to be available first quarter 2011, with additional PTFs. All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

© 2010 IBM Corporation

z/OS Crypto Toleration Support (1 of 2)

z/OS Release	Crypto Web Download	FMID	Comments
z/OS V1.7[1]	N/A – In z/OS base product	HCR7720	▪NO Toleration Support on z196
	[09/2005] Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7 (web download no longer available)	HCR7730	▪Enhancements to CPACF (AES, PRNG, and SHA-256) and Crypto Express2 exploitation ▪NO Toleration Support on z196
	[09/2005] Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/R7 (web download no longer available)	HCR7731	▪CPACF and Crypto Express2 exploitation ▪Enhancements include Remote Key Loading for ATMs and POSs
	[11/2007] Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 web deliverable	HCR7750	▪4096-bit RSA keys, ISO-3 PIN Support, CPACF support for SHA-394 and SHA-512, Reduced support of retained private key in ICSF
	[11/2008] Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8 web deliverable	HCR7751	▪Support for 13- thru 19-digit Personal Account Numbers ▪ ICSF Query Algorithms, ▪ Key Token Policy support ▪Secure Key AES

[1] The IBM Lifecycle Extension for z/OS is required for support after general support is withdrawn

© 2010 IBM Corporation

z/OS Crypto Toleration Support (2 of 2)

z/OS Release	**Crypto Web Download**	**FMID**	**Comments**
z/OS V1.8[1]	N/A – In z/OS base product	HCR7731	▪CPACF and Crypto Express2 exploitation ▪Enhancements include Remote Key Loading for ATMs and POSs
	[11/2007] Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 web deliverable	HCR7750	▪4096-bit RSA keys, ISO-3 PIN Support, CPACF support for SHA-394 and SHA-512, Reduced support of retained private key in ICSF
	[11/2008] Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8 web deliverable	HCR7751	▪ Support for 13- thru 19-digit Personal Account Numbers ▪ ICSF Query Algorithms, ▪ Key Token Policy support ▪ Secure Key AES
z/OS V1.9[1]	N/A – In z/OS base product	HCR7740	▪ PKCS11 Support
	[11/2007] Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 web deliverable	HCR7750	▪ 4096-bit RSA keys, ISO-3 PIN Support, CPACF support for SHA-394 and SHA-512, Reduced support of retained private key in ICSF
	[11/2008] Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8 web deliverable	HCR7751	▪ Support for 13- thru 19-digit Personal Account Numbers ▪ ICSF Query Algorithms, ▪ Key Token Policy support ▪ Secure Key AES
	[11/2009] Cryptographic Support for z/OS V1R9-V1R11 Web deliverable	HCR7770	▪ Protected Key CP Assist for Cryptographic Function ▪ New Crypto Express3 and Crypto Express3 -1P

[1] The IBM Lifecycle Extension for z/OS is required for support after general support is withdrawn

© 2010 IBM Corporation

z/OS Crypto Toleration and Exploitation Support

z/OS Release	Crypto Web Download	FMID	Comments
z/OS V1.10	N/A – In z/OS base product	HCR7750	▪All of the above
	[11/2008] Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8 web deliverable	HCR7751	▪Support for 13- thru 19-digit Personal Account Numbers ▪ ICSF Query Algorithms, ▪ Key Token Policy support ▪Secure Key AES
	[11/2009] Cryptographic Support for z/OS V1R9-V1R11 Web deliverable	HCR7770	▪Protected Key CP Assist for Cryptographic Function ▪New Crypto Express3 and Crypto Express3 -1P
	[planned 9/2010] Cryptographic Support for z/OS V1R10-V1R12 Web deliverable	HCR7780	▪X9.8 Pin, 64 Bit, z196 CPACF, HMAC*, CKDS Constraint Relief, AP Interrupt, PCI Audit, ECC HW Support, CBC Key Wrap
z/OS V1.11	N/A – In z/OS base product	HCR7751	▪All of the above
	[11/2009] Cryptographic Support for z/OS V1R9-V1R11 Web deliverable	HCR7770	▪Protected Key CP Assist for Cryptographic Function ▪New Crypto Express3 and Crypto Express3 -1P
	[planned 9/2010] Cryptographic Support for z/OS V1R10-V1R12 Web deliverable	HCR7780	▪X9.8 Pin, 64 Bit, z196 CPACF, HMAC*, CKDS Constraint Relief, AP Interrupt, PCI Audit, ECC HW Support, CBC Key Wrap
z/OS V1.12	N/A – In z/OS base product	HCR7770	▪All of the above
	[planned 9/2010] Cryptographic Support for z/OS V1R10-V1R12 Web deliverable	HCR7780	▪X9.8 Pin, 64 Bit, z196 CPACF, HMAC*, CKDS Constraint Relief, AP Interrupt, PCI Audit, ECC HW Support, CBC Key Wrap

* Planned to be available first quarter 2011, with additional PTFs. All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

© 2010 IBM Corporation

External Name	FMID	Applicable z/OS Releases	Avail	EoM
z/OS or z/OS.e V1.3 or V1.4	HCR7706	z/OS V1.3 and z/OS V1.4	3/2002	9/2004*
z990 Cryptographic CP Assist Support for z/OS V1.3[1]	HCR7708	z/OS V1.3	6/2003	10/2003
z/OS V1.4 z990 Compatibility Support or z/OS.e z990 Coexistence	HCR7708	z/OS V1.4	6/2003	10/2003
z/OS V1.4 z990 Exploitation Support or z/OS.e Coexistence Update feature	HCR7708	z/OS V1.4	10/2003	12/2006
z990 Cryptographic Support [2]	HCR770A	OS/390 V2.10, z/OS 1.2, z/OS 1.3, z/OS V1.4, any z/OS V1.4 features, and z/OS V1.5	9/2003	5/2004
z/OS or z/OS.e V1.5	HCR7708	z/OS V1.5	3/2004	9/2004
z/OS or z/OS.e V1.6	HCR770A	z/OS V1.6	9/2004	10/2005*
z990 and z890 Enhancements to Cryptographic Support [3]	HCR770B	OS/390 V2.10, z/OS V1.2, z/OS V1.3, z/OS V1.4 and z/OS V1.5	5/2004	12/2006
ICSF 64-bit Virtual Support for Z/OS V1.6 and z/OS.e V1.6 [4]	HCR7720	z/OS V1.6	12/2004	9/2005
z/OS or z/OS.e V1.7	HCR7720	z/OS V1.7	9/2005	10/2006
Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7[5]	HCR7730	z/OS V1.6 and z/OS V1.7	9/2005	5/2006
Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/R7[6]	HCR7731	z/OS V1.6 and z/OS V1.7	5/2006	11/2007
z/OS or z/OS.e V1.8	HCR7731	z/OS V1.8	9/2006	10/2007
z/OS V1.9	HCR7740	z/OS V1.9	9/2007	10/2008
Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 web deliverable[7]	HCR7750	z/OS V1.7, z/OS V1.8 and z/OS V1.9	9/2007	11/2008
z/OS V1.10	HCR7750	z/OS V1.10	9/2008	10/2009
Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8 web deliverable[8]	HCR7751	z/OS V1.7*, z/OS V1.8, z/OS V1.9, z/OS V1.10	11/2008	11/2009
z/OS V1.11	HCR7751	z/OS V1.11	9/2009	Planned 10/2010
Cryptographic Support for z/OS V1R9-V1R11 Web deliverable[9]	HCR7770	z/OS V1.9, z/OS V1.10, z/OS V1.11	11/2009	Planned 9/2010
z/OS V1.12	HCR7770	z/OS V1.12	9/2010	Planned 10/2011
Cryptographic Support for z/OS V1R10-V1R12 Web deliverable[10]	HCR7780	z/OS V1.10, z/OS V1.11, z/OS V1.12	9/2010	TBD

© 2010 IBM Corporation

STP Planning

© 2010 IBM Corporation

STP References and Additional Information

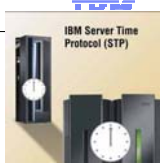

- **Redbooks**
  - Server Time Protocol Planning Guide, SG24-7280
    - http://www.redbooks.**ibm.com**/redpieces/abstracts/sg247280.html
  - Server Time Protocol Implementation Guide, SG24-7281
    - http://www.redbooks.**ibm.com**/redpieces/abstracts/sg247281.html
- **Education**
  - Introduction to Server Time Protocol (STP)
    - Available on Resource Link at General Availability (GA)
    - www.**ibm.com**/servers/resourcelink/hom03010.nsf?OpenDatabase
- **STP Web site**
  - www.**ibm.com**/systems/z/pso/stp.html
- **Systems Assurance**
  - The IBM team is required to complete a Systems Assurance Review (SAPR Guide, SA06-012) and to complete the Systems Assurance Confirmation Form via Resource Link
  - http://w3-03.**ibm.com**/support/assure/assur30i.nsf/WebIndex/SA779

© 2010 IBM Corporation

CF Link Connectivity

© 2010 IBM Corporation

System z9 Link Connectivity to z9, z10, z196

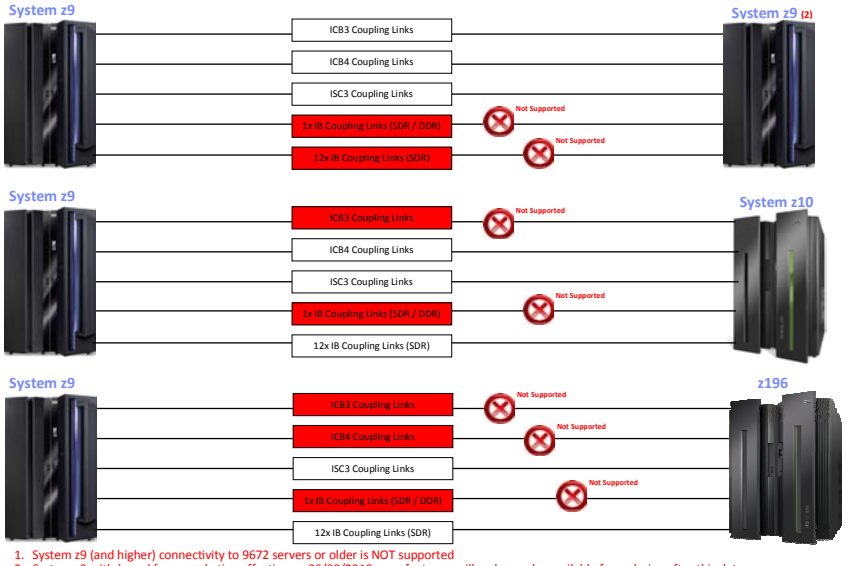

1. System z9 (and higher) connectivity to 9672 servers or older is NOT supported
2. System z9 withdrawal from marketing effective on 06/30/2010, new feature s will no longer be available for ordering after this date

© 2010 IBM Corporation

System z10 Link Connectivity to z9, z10, z196

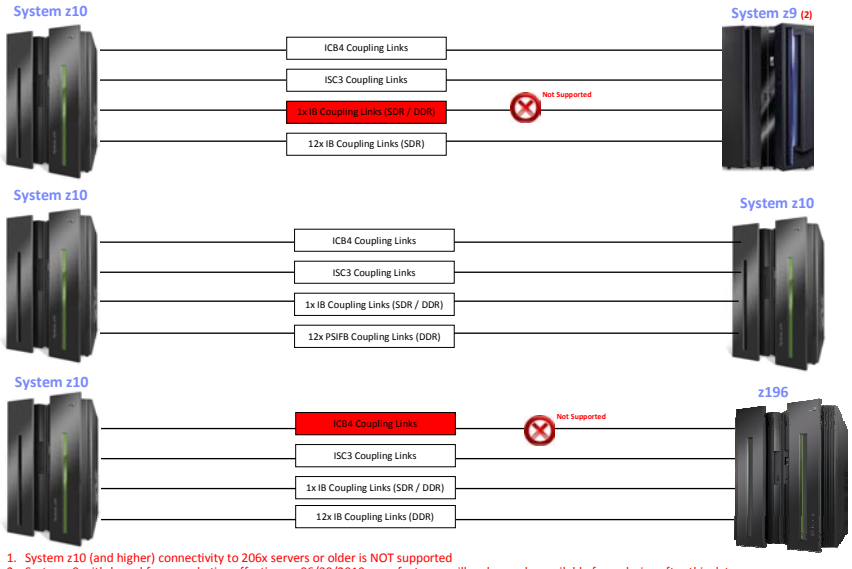

1. System z10 (and higher) connectivity to 206x servers or older is NOT supported
2. System z9 withdrawal from marketing effective on 06/30/2010, new feature s will no longer be available for ordering after this date

© 2010 IBM Corporation

System z196 Link Connectivity to z9, z10, z196

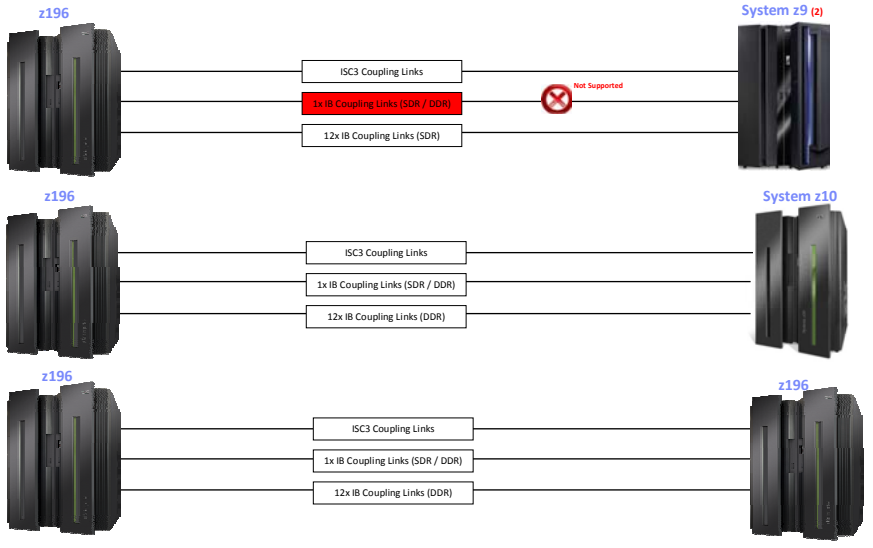

1. z196 connectivity to 208x servers or older is NOT supported
2. System z9 withdrawal from marketing effective on 06/30/2010, new feature s will no longer be available for ordering after this date

© 2010 IBM Corporation

© 2010 IBM Corporation

zEnd

The Future Runs on System z

© 2010 IBM Corporation