

Slide 1



What's New in z/OS, Beantown Edition

Session 2224




John Eells
IBM Poughkeepsie
eells@us.ibm.com
2 August 2010

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

Slide 2



Trademarks

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

AIX*	DS6000	Infiniband*	ServerPac*	z10
BladeCenter*	DS8000*	InfoSphere	SystemPac*	z10 BC
CICS*	FICON*	Language Environment*	System Storage	z10 EC
DB2*	FlashCopy*	MVS	System z	z/OS*
DFSMS	Hiperspace	Parallel Sysplex*	System z9	zEnterprise
DFSMSdtp	IBM*	ProductPac*	System z10	zSeries*
DFSMSdas	IBM eServer	RACF*	System z10 Business Class	
DFSMSnam	IBM logo*	Redbooks*	Truolt*	
DFSMSsmm	ibm.com	REXX	WebSphere*	
DFSORT	IMS	RMF	z9*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.


InfiniBand is a registered trademark of the InfiniBand Trade Association (IBTA).
Intel is a trademark of the Intel Corporation in the United States and other countries.
Linux is a trademark of Linux. Trademarks in the United States, other countries, or both.
Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
UNIX is a registered trademark of The Open Group in the United States and other countries.
All other products may be trademarks or registered trademarks of their respective companies.
The Open Group is a registered trademark of The Open Group in the US and other countries.

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 multiprocessing 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.
This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

2

©2010 IBM Corporation


Slide 3




A zNextGen Bonus!

- Mastering IBM System z®—Entry Level System Programmer Certification
- Free test:


Monday, 6:00-7:30 p.m., Room 311, Hynes Convention Center



Slide 4



z/OS Support Summary



■ Out of service

■ Lifecycle Extension withdrawal 2 years later

Service Withdrawal Dates

z/OS®	z800/ z900	z890/ z990	z9® EC	z9 BC	z10 EC™	z10 BC™	z196	DS8000® DS6000™	TS1130	End of Service	Coexists with z/OS...	Planned Ship Date²
R6	x	x	x	x				x		9/07		R8
R7	x	x	x	x	x⁴	x³	x⁴	x⁴	x	9/08¹		R9
R8	x	x	x	x	x	x	x	x	x	9/09¹		R10
R9	x	x	x	x	x	x	x	x	x	9/10¹		R11
R10	x	x	x	x	x	x	x	x	x	9/11		R12²
R11	x	x	x	x	x	x	x	x	x	9/12²		R13²
R12	x	x	x	x	x	x	x	x	x	9/13²		R14²
R13²	x	x	x	x	x	x	x	x	x	9/14²		R16² 9/11²


Migrating to z/OS 1.12 Parts 1 and 2
(Wednesday 3:00, Wednesday 6:00)

1. Fee-based service extension available

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

3. IBM Lifecycle Extension for z/OS V1.7 (6637-A01) required

4. Fee-based service extension required for some features



This chart summarizes much of the information about z/OS that you need to know. Server and DASD hardware support, end of service dates, coexistence, and planned availability dates are shown here for existing and planned releases through z/OS R13.

Note that:

The lowest supported level of z/OS is R9 (without Lifecycle Extension)
R11 is the currently orderable release (until September 2010)

Hardware Support

z/OS and IBM zEnterprise Functions and Features¹

Five hardware models
Increased capacity processors
Up to 15 subcapacity CPs at capacity settings 4, 5, or 6
Up to 3 TB RAIM (real) memory ²
6.0 GB/sec InfiniBand [®] I/O interconnect
8 slot, 2 domain I/O drawer
Concurrent I/O drawer add, remove, replace
Optional water cooling
Optional High Voltage DC power
Optional overhead I/O cable exit
Up to 80 processors per server configurable as CPs, zAAPs, zIIPs, IFLs, ICFs, or SAPs (up to 32-way on R7, 64-way on R9, 80-way on R11)
New and enhanced instructions

Capacity Provisioning enhanced ⁴
Three subchannel sets per LCSS ³
Platform Management from HMC
CFCC Level 17 enhancements ⁴
Up to 128 Coupling Link CHPIDs
Improved processor cache design
Power save functions
Crypto Express3 enhancements ⁵
Secure key HMAC Support
Elliptic Curve Cryptography (ECC) Digital Signatures ³
CPACF enhancements ⁵
Out of order instruction execution
z/OS discovery and auto-configuration (zDAC) ³
OSA-Express3 Inbound Workload Queuing (IWQ) ³

z/OS exploitation in blue

IBM zEnterprise™ Overview: Tuesday 11:00

z/OS Software Positioning for the Latest System z Hardware: Thursday 3:00

1. z/OS R7 and z/OS R8 support requires IBM Lifecycle Extension for z/OS (5637-A01 or 5638-A01). PTFs required for z/OS R7-R12; refer to the PSP.

2. Maximum of 1 TB per LPAR. Maximum supported by z/OS R7 is 512 GB. z/OS R8 and later are designed to support up to 4 TB per image.

3. z/OS R12 required

4. z/OS R12, or R10 or later with PTFs required

5. 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, required*

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

Integration and exploitation of IBM zEnterprise System

IBM introduces the IBM zEnterprise System -- a system that combines the gold standard of enterprise computing with built-in function to extend IBM's mainframe-like governance and qualities of service to special-purpose workload optimizers and general-purpose application serving. End-to-end management is enabled for this heterogeneous environment by the IBM zEnterprise Unified Resource Manager, which provides energy monitoring and management, goal-oriented policy-based workload monitoring and management, increased security, virtual networking, and data management, consolidated in a single interface that can be tied to business requirements. An IBM zEnterprise System is composed of the IBM zEnterprise 196, the IBM zEnterprise Unified Resource Manager, the IBM zEnterprise BladeCenter Extension (zBX), and optimizers or blades. 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 (2817-780) 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, and is supported by z/OS V1.12. This support is also available for z/OS V1.10 and z/OS V1.11 with PTFs.

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. z/OS V1.12 on zEnterprise 196 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. This support requires the PTF for APAR OA32807; PTFs are also available for z/OS V1.10 and z/OS V1.11.

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.

A new IBM System z Discovery and Auto-Configuration (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.

z/OS V1.12 on z196 servers is designed to provide improved networking performance with OSA-Express-3 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 Express-3 IWQ and its ability to reduce the amount of costly network retransmissions (by reducing the incidence of out-of-order packets). A streamlined CommServer execution path for Sysplex Distributor over IWQ is expected to improve performance for sysplex-distributed traffic as well.

When inbound workload queueing 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 queueing 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.

IWQ is supported on z/OS V1.12 and is exclusive to OSA-Express3 on z196 and System z10, where CHPID types OSM and OSX are exclusive to z196.

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.

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, BladeCenters, blades, top-of-rack switches, power distribution units, and hypervisors) for management purposes. This INMN is pre-wired, internally switched, configured, and managed with full redundancy for high availability. These characteristics of the INMN enable simplified and secure management with no effect on your application data traffic. The INMN is exclusive to zEnterprise System, is supported by z/OS V1.12, 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, is supported by z/OS V1.12, and is configured with a CHPID type of OSX.

z/OS V1.12 Communications Server adds support for OSA-Express3 adapters configured with the new OSM and OSX CHPID types, thus allowing z/OS TCP/IP connectivity to these two internal networks. 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 Intra-Ensemble Data Network.

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), in z/OS V1.12 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.

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. These options are 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 compute-intensive integer workload 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.

Slide 6

Hardware Support

zEnterprise and z/OS Support for New Functions

	Minimum z/OS Level *
Basic zEnterprise server support	1.7 ^{1,2,4,6}
Tolerance for more than 64 CPs on a server	1.7 ^{1,2,4}
Crypto Tolerance	1.7 ^{1,2,4,6}
HiperDispatch node and cache support	1.10 ^{1,2}
Power Save Mode	1.10 ^{1,2}
CPU Measurement Facility	1.10 ^{1,2}
CF Level 17 exploitation	1.10 ^{1,2,7}
zHPF Response Time Reduction	1.10 ^{1,2}
Three Subchannel Sets	1.10 ^{1,2}
Crypto Exploitation	1.10 ^{1,2,3}
RMF™ Postprocessor Support for 4096-bit Crypto	1.10 ^{1,2}
IBM zEnterprise BladeCenter® Extension (zBX) Support, including Network Management and Performance Management	1.10 ^{1,2}
80-Way Single System Image Support	1.11 ¹
XL C/C++ ARCH(9) TUNE(9)	1.12
XL C/C++ HW Exploitation	1.12
zDAC Support	1.12
D OSAINFO Command Support	1.12
OSA-Express3 Inbound Workload Queuing	1.12

Hardware Overview
Tuesday 1:30

1. Additional features, service, or Web downloads required. 2. Please refer to the current PSP buckets for latest PTFs for IBM zEnterprise 106 (z196) Compatibility and new functions/features support. 3. z/OS V1.10, z/OS V1.11, or z/OS V1.12 with the Cryptographic Support for z/OS V1R10-V1R12 Web deliverable. 4. Lifecycle Extension required. 5. Crypto web deliverable required. 6. z/OS support web deliverable required. HiperDispatch support on z/OS 1.7.7. Some functions require z/OS V1.12.

* 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

Slide 7

z/OS 1.12

A smarter operating system

Improving Usability and Skills
zDAC, Multiple TSO/E logons in a JESplex, METAL C headers for Health Checker and sample check, Health check history display, More checks, z/OSMF WLM Policy Editor and monitoring, Incident Log improvements, DFSORT™, ...

Integrating new Applications and Supporting Industry and Open Standards
XML dynamic schema and fragment parsing; Language Environment® support for calendar dates beyond 2038, C/C++ enhancements, and VSAM AIX in EAS on EAV; z/OS UNIX® support for FILEDATA=RECORD, ...

Scalability & Performance
80-way single image, Three subchannel sets, more EAV support, Larger catalogs, OAM threadsafe, Parallel Sysplex® scalability improvements, Large Pages for NUCLEUS, OSA-Express3 Inbound Workload Queuing (IWQ), XML parsing, XL C/C++, ...

Enhancing Security
Support for Crypto Express3 and CPACF enhancements, VISA smart card support, LDAP ACL by IP address and password policy; ECC support in RACF®, System SSL, and PKI; Health Checker support for BPX.SUPERUSER, SAF controls on sysplex-scope CPF commands, ...

Improving Availability
VSAM CA Reclaim; Faster DB2® startup; XCF monitoring of critical members; SVC Dump improvements; SMF record flood detection and automation, filtering; PFA extensions; Run Time Diagnostics, Concurrent CF dump, ...

Self Managing Capabilities
HiperDispatch updates for z196, Capacity Provisioning interval averaging, CP monitoring of CICS® and IMS™ transaction classes, Improved WLM batch management, WTOR Auto-Reply, ...

Extending the Network
Continued IPv6, security related RFC focus; TN3270 VSCR; Trusted Connections, AT-TLS Improvements; IKEv2, Sysplex Distributor Hot Standby, ...

* 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

At a glance

This latest release, z/OS® V1.12, delivers truly significant improvements in system availability, workload performance, simplified usability, and cross system integrated connectivity. z/OS V1.12 takes smart systems to a whole new dimension by providing automatic and real-time capabilities for higher performance and less operator

intervention, with fewer system disruptions and response time impacts to z/OS and the business applications that rely on z/OS.

Avoiding data fragmentation and data reorganizations

With new VSAM Control Area (CA) Reclaim capability, applications that use VSAM key-sequenced data sets (KSDS) can benefit from improved availability through the avoidance of outages that used to be required to defragment and reorganize this data, as well as improved performance and space utilization.

Real-time decision making in the event of a system problem

A new z/OS Runtime Diagnostics function is designed to quickly look at the system message log and address space resources and can help you identify sources of possible problems in as little as one minute, giving your experienced operators the information they need to act swiftly.

Predicting problems

z/OS Predictive Failure Analysis® (PFA) is extended to monitor the rate at which SMF records are generated. When the rate is abnormally high the system is designed to issue an alert to warn you, giving you time to proactively react prior to a system issue.

Performance for many key workloads

Examples of performance improvements include: up to 44% for VSAM-based workloads, shorter DB2® 9 for z/OS restart time and 50-90% for SVC dump capture time workloads (see Note).

New XML fragment validation

By revalidating only the fragments being updated, DB2 10 for z/OS pureXML can avoid the costly revalidation of entire XML documents, which without this function can take many times longer.

Networking performance improvements

Networking throughput for interactive workloads can be improved by 30-50% and secure Application Transparent, Transport Layer Security (AT-TLS) can be improved by up to 30% with a reduction in CPU overhead (see Note).

Improved productivity with the new face of z/OS

The z/OS Management Facility (5655-S28) can help shave hours of system administration time down to minutes and minutes of time down to seconds, and can enable new system programmers to be productive in weeks.

Fast FICON® configuration

New z/OS FICON Discovery and Auto Configuration (zDAC) can automatically discover new and updated FICON fabric-attached devices and propose channel configurations based on high availability best practices and existing configurations even taking into account Parallel Sysplex® symmetric configurations. This function is sophisticated and

takes advantage of the robust I/O subsystem that is available with the new IBM® zEnterprise™ server.

Take advantage of a new class of optimizers

DB2 9 for z/OS is the first exploiter of the IBM zEnterprise BladeCenter® optimizers. 5697-AQT, IBM Smart Analytics Optimizer DB2 for z/OS V1.1, announced in Software Announcement 210-266, dated July 22, 2010, is designed to help mitigate database administration costs while speeding up the performance of typical data warehouse or OLAP-specific workloads by an order of magnitude (see Note).

Note: Based on IBM Lab results; your results will vary.

- VSAM performance improvement is through the use of VSAM CA Reclaim; actual benefit may be more or less and will depend on the degree of VSAM data fragmentation and how the data is accessed. It is anticipated that VSAM key sequenced data sets (KSDs) that are severely fragmented or rarely reorganized will see the most benefit. For applications that delete a large number of records from a narrow key range and then immediately re-insert them, CA Reclaim could result in some performance degradation.
- DB2 9 for z/OS startup time reduction is through z/OS Allocation, DFSMSdfp™, and GRS improvements; actual benefit will depend on the number of data sets opened. It is anticipated that address spaces opening up many thousands of data sets will see more benefit.
- Actual SVC dump time will depend on the amount of data being captured and the amount of that data dumped from auxiliary storage.
- The AT-TLS CPU consumption results were obtained on System z10®, model 2097-E64. Actual AT-TLS CPU consumption improvement will depend on the amount of data being transmitted and whether the workload is interactive or streaming. Throughput gain due to this improvement in CPU consumption is likely, but would vary depending on overall utilization of the z/OS image.
- Interactive throughput measurements were obtained on System z10, model 2097-E64 with OSA-Express3 Inbound Workload Queuing function. Actual benefits will depend on the 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.
- Actual IBM Smart Analytics Optimizer CPU-intensive query performance improvement will depend on the specific query workload.

Overview

z/OS V1.12 provides significant value to your business by delivering advanced real-time capabilities designed to provide higher performance, improved productivity, new levels of data availability, and fewer disruptions and impacts to z/OS and the business applications that rely on z/OS.

Performance and availability

This release of z/OS delivers real-time capabilities for application performance and data availability. With the new VSAM Control Area (CA) Reclaim capability, applications that use VSAM key-sequenced data sets (KSDS) can benefit from improved performance, minimized space utilization, and improved application availability. For example, IBM system tests indicate performance of dedicated VSAM workload could improve by up to 44% (see Note), in addition to helping you avoid outages that used to be required to defragment and reorganize this data. Performance improvements are anticipated for many applications using CICS®, VSAM, VSAM RLS, IMS™ VSAM, and Catalog processing.

There are significant performance enhancements expected for other workloads as well. For example, DB2 9 for z/OS (5635-DB2) startup time could be reduced (see Note), and z/OS XML System Services validating parsing performance is anticipated to improve by 30-50% (see Note). New fragment validation parsing helps DB2 X for z/OS (5605-DB2) pureXML avoid costly revalidation of the whole XML document. In addition, SVC dump capture time can be reduced by up to 50-90% (see Note).

Network performance is improved. Processing overhead for Application Transparent - Transport Layer Security (AT-TLS) can be improved by 30% (see Note). Network throughput for interactive workloads can be improved by 30-50% (see Note) using the new Inbound Workload Queuing (IWQ) function, which is exclusive to OSA-Express3 on IBM zEnterprise 196 (z196) and System z10 servers. It is anticipated that IWQ will provide benefit for streaming workloads by reducing the number of costly network retransmissions due to out-of-order packets, and will benefit Sysplex Distributor traffic by providing a more streamlined processing path.

Smart technologies

Predictive Failure Analysis (PFA) has been updated to monitor SMF data rates for z/OS V1.12.

In the rare event of a system issue, a new base function, z/OS Runtime Diagnostics, is designed to quickly analyze key indicators on a running system and help identify the root causes of system problems and degradations. The Runtime Diagnostics function is anticipated to run in as little as one minute to return results fast enough to help you choose between alternative corrective actions swiftly and effectively and facilitate high levels of system and application availability.

The new z/OS

Time for a new look at z/OS. The z/OS Management Facility (5655-S28) is the new face of z/OS. It can shave hours of time from system management tasks, and with embedded guidance and best practices built in, new system programmers can be productive in weeks. Capturing and analyzing incident data, configuring TCP/IP network settings and z/OS Workload Management (WLM) policies, and monitoring z/OS performance data have never been easier. z/OS Management Facility V1.12 is in "lock step" with z/OS V1.12 and is ready to help you use many of its latest features.

z/OS V1.12 provides simplified configuration of FICON-connected disk and tape through the revolutionary new z/OS FICON Discovery and Auto Configuration (zDAC) capability. zDAC can automatically discover new or updated devices and propose channel configurations, helping preserve Parallel Sysplex symmetric configurations. This function is sophisticated and takes advantage of the robust I/O subsystem that is available only with the new IBM zEnterprise 196 server.

Connected and ready

IBM introduces the new IBM zEnterprise System, a revolutionary direction in multi-tier, multi-system, multi-workload management and coordination. The zEnterprise System is beneficial for all System z® workloads, but is especially ideal for z/OS workloads because it supports and strengthens your z/OS environment by targeting the technology and application tiers that, in turn, rely on these core systems.

The combination of z/OS and the zEnterprise System can provide workload management integration and insight (through integration with the zEnterprise Unified Resource Manager), designed to provide both high speed and highly secure networking (through the use of the physically isolated data and service networks), and infrastructure simplification (with fewer hops, devices, servers, and processors that can complicate today's data centers).

z/OS is designed to integrate seamlessly with the zEnterprise System -- only a handful of new configuration settings are needed to integrate and execute in the new environment. In addition, DB2 9 for z/OS and the IBM Smart Analytics Optimizer for DB2 for z/OS V1.1 (5697-AQT) are the first exploiters of the IBM zEnterprise BladeCenter optimizers. The IBM Smart Analytics Optimizer is designed to help mitigate database administration costs while speeding up the performance of typical data warehouse or OLAP-specific workloads by an order of magnitude (see Note).

Get the best of all worlds with z/OS V1.12 and the new IBM zEnterprise System. Together, this revolutionary platform can extend the strengths of z/OS by deploying multi-tier workloads in a consolidated, simplified, and highly integrated environment up close to z/OS.

Note: Based on IBM Lab results; your results will vary.

- VSAM performance improvement is through the use of VSAM CA Reclaim; actual benefit may be more or less and will depend on the degree of VSAM data fragmentation and how the data is accessed. It is anticipated that VSAM key sequenced data sets (KSDS) that are severely fragmented or rarely reorganized will see the most benefit. For applications that delete a large number of records from a narrow key range and then immediately re-insert them, CA Reclaim could result in some performance degradation.
- DB2 9 for z/OS startup time reduction is through z/OS Allocation, DFSMSdfp, and GRS improvements; actual benefit will depend on the number of data sets opened. It is anticipated that address spaces opening up many thousands of data sets will see more benefit.

- z/OS XML System Services validating parsing performance improvement will depend on the amount of data being parsed and the degree of complexity of the schema.
- Actual SVC dump time will depend on the amount of data being captured and the amount of that data dumped from auxiliary storage.
- The AT-TLS CPU consumption results were obtained on System z10, model 2097-E64. Actual AT-TLS CPU consumption improvement will depend on the amount of data being transmitted and whether the workload is interactive or streaming. Throughput gain due to this improvement in CPU consumption is likely, but would vary depending on overall utilization of the z/OS image.
- Interactive throughput measurements were obtained on System z10, model 2097-E64 with OSA-Express3 Inbound Workload Queuing function. Actual benefits will depend on the 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.
- Actual IBM Smart Analytics Optimizer CPU-intensive query performance improvement will depend on the specific query workload.

Key prerequisites

z/OS V1.12 runs on these IBM System z servers:

- z196
- z10™ EC
- z10 BC
- z9™ EC (see Note)
- z9 BC (see Note)
- z990 (see Note)
- z890 (see Note)
- z900 (see Note)
- z800 (see Note)

Note: These products are withdrawn from marketing.

For a complete description of z/OS V1.12 hardware prerequisites, refer to *z/OS VIR12 Planning for Installation* (GA22-7504), when available, at http://publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/BOOKS/E0Z2B1B0

Planned availability date

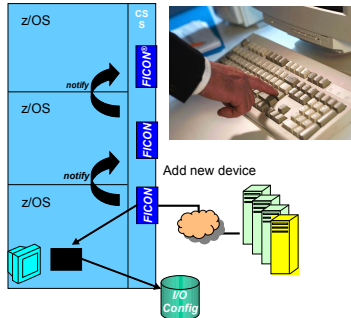
September 24, 2010

Improving Usability and Skills

Discovery and Autoconfiguration for DASD and Tape (zDAC)*

- Set up a policy for what to do when new devices are detected (one-time)
- Configure new storage (e.g., with TPC) and attach it to the system
- Use a new HCD option to process newly-discovered devices
- HCD creates a configuration with a device number range, base and dynamic CHP assignments, DCM, PAV and HyperPAV, etc.
- Once you approve, a work IODF is created with IODF, LPAR, switch, and control unit configuration changes
- Finally, issue an ACTIVATE command

***Note: Control units must register with the name server to be discoverable.**



IOS530I DISCOVERY AND AUTOCONFIG PROCESSING WAS SUCCESSFUL...

IOS502I I/O CONFIGURATION CHANGED
NEW IODF = newiodf ...

8
© 2010 IBM Corporation

Ease of use

IBM has made a huge commitment to simplifying the use, operation, and administration of mainframe systems. The past several releases of z/OS delivered improvements in the areas of simplifying diagnosis and problem determination, network and security management, overall z/OS I/O configuration, and sysplex operations. New automatic, real-time capabilities can mean less operator intervention, and built-in expert guidance can help reduce time to perform tasks. These improvements are designed to help simplify systems management, to improve application programmer, system programmer, and operator productivity, and to provide fewer opportunities for the introduction of human errors.

The introduction of z/OS Management Facility (z/OSMF) was a significant step. The z/OS Management Facility is the new face for z/OS and it supports a modern, Web-browser-based management console for z/OS. Automated tasks can help reduce the learning curve and improve productivity for system programmers. For example, new users may quickly become proficient in performing tasks using z/OSMF applications. In addition, embedded active user assistance, such as wizards, can guide you through tasks and provide simplified operations, for example, reducing hours of tasks down to minutes, or minutes worth of tasks down to seconds. It is the infrastructure from which valuable new applications and enhancements will be made.

IBM's commitment to mainframe simplification is not isolated to z/OS. It is integrated into the platform stack as well. CICS Explorer™ provides CICS architects, developers, system programmers, and administrators a common tooling environment with integrated access to a wide range of data and control capabilities. CICS Deployment Assistant for z/OS V1.1, announced in Software Announcement 210-221, dated July 22, 2010, further simplifies CICS by adding discovery, visualization, and automated deployment of CICS

regions, including the ability to edit and submit z/OS JCL directly from the CICS Explorer. DB2 Data Studio provides an integrated set of tooling to support all phases of the data management life cycle. IMS provides a new integrated development environment and operational console to accelerate the development time for new IMS (5635-A02) applications and optimize collaboration. Rational® Developer for System z (5724-T07) helps simplify collaboration, development, and delivery of business applications and integrate existing core business applications with Web services and SOA. Tivoli® Service Management Center (5724-M19) provides a set of integrated solutions and building-blocks that allows a business to implement an enterprise-wide service management and process automation hub on System z. These technologies have the power to reduce application development, deployment, and management times significantly.

Simplification enhancements specific to z/OS V1.12 are:

The IBM z/OS Management Facility V1.12, in conjunction with z/OS V1.12, adds two new system management tasks: Workload Management (policy editor), and Sysplex Status and Monitoring Desktops (for monitoring resources). z/OSMF V1.12 also provides enhancements to the already valuable Incident Log and Configuration Assistant for the z/OS Communications Server functions, and the ability to add other application launch points and links to the z/OSMF menu. The Configuration Assistant for the z/OS Communications Server has been updated with added support for IP security: IKEv2, new cryptographic and authentication algorithms, and enhancements to the support for Application Transparent Transport Layer Security (AT-TLS). The z/OS Management Facility Workload Management (policy editor) task can help reduce many hours of manual tasks to minutes. The Sysplex Status and Monitoring Desktops tasks can provide real-time performance status of all your z/OS servers and sysplexes, as well as Linux images from one location, and it takes just a few key clicks to view real-time performance metrics. The Incident Log adds support for encrypted parallel FTP dumps to be sent to IBM and allows you to add additional custom information to individual incidents.

A new IBM System z Discovery and Auto-Configuration (zDAC) function, available with z/OS V1.12 on the 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. 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.

IBM Health Checker for z/OS has a long history of helping to simplify and automate the identification of potential configuration problems before they impact system availability by comparing active values and settings to those suggested by IBM or defined by your installation. The z/OS Health Checker is extremely valuable not only in identifying exceptions to z/OS configurations, but also in identifying migration actions and checking that these migration actions are completed accurately. In addition, reports from the z/OS Health Checker may be used to support your corporate compliance. For example, z/OS Health Checker reports can help identify unsecured resources that should be RACF protected, can help validate the redundancy in a Parallel Sysplex configuration, and could be used as part of risk assessment exercises. For z/OS V1.12, Health Checker is updated with the ability to write checks in Metal C, and with the addition of more checks for Parallel Sysplex, such as best practices for coupling facility structure size, couple data set specification limits, Sysplex Failure Manager policies, and coupling facility allocation. Also, there are new checks for SMB server, DFSMS™, I/O Supervisor, TCP/IP IPv4 and IPv6 usage, HFS to zFS migration, and still others.

There are additional ease of use enhancements designed to help prevent many JCL errors arising from duplicate temporary data set names, simplify Language Environment® and zFS migration, simplify RMF processing, improve performance, and allow users to create customized views for the Library Server.

Details on the ease of use and platform simplification enhancements included in z/OS V1.12:

- The following functions are included in z/OSMF V1.12:
 - The industry's workload management gold standard, z/OS Workload Manager (z/OS WLM), just got easier to use with a new z/OSMF function that makes it easier to create, modify, and review WLM service definitions. With the new z/OS Management Facility Workload Management (policy editor) systems management task, many hours of manual steps can be reduced to minutes. For example, while optimizing a service definition based on best practices, hours of reading through WLM-related manuals are reduced to minutes due to the z/OSMF built-in best practices.
 - The Sysplex Status and Monitoring Desktops systems management tasks can provide real-time status of all your servers, sysplexes, and Linux images from one location. Sysplex Status can provide cross-sysplex performance monitoring from a single point of control with quick red-yellow-green system health indicators on a single panel so it takes mere seconds to check all your sysplexes, as well as Linux for System z, and Linux for Intel® images. For example, if workload in one sysplex is missing goals, you can quickly drill down to system components using the Monitoring Desktops application where it can take as few as three keystrokes to view real-time statistics.
 - z/OSMF V1.12 has been enhanced so that links and launch points to non-z/OSMF resources can be added to any category in the Navigation tree on

the left-hand side of the screen and not just the Links category. A new script interface enables you to programmatically add your own non-z/OSMF launch points and links to the left-hand side navigation tree so system programmers have the ability to access more information and tools quickly.

- A large portion of the z/OS Management Facility application is written in Java™, and therefore is eligible to run on IBM System z Application Assist Processors (zAAPs). Some functions in z/OS Management Facility use the Common Information Model (CIM) and, starting with z/OS V1.11, are eligible for the IBM System z Integrated information Processor (zIIP). Also starting with z/OS V1.11, zAAP-eligible workloads may run on the zIIP processor, thus helping maximize the return on your zIIP investment. For the new zAAP on zIIP capability, additional terms and conditions apply. See 5694-A01, z/OS V1.11 - Smart, adaptive, trusted, efficient, announced in Software Announcement 209-242, dated August 18, 2009.
- z/OSMF V1.12 supports browsers running on Microsoft Windows XP, Windows Vista, and Windows 7 Professional Edition (32-bit). Supported browsers are Microsoft Internet Explorer 7 and 8, as well as Mozilla Firefox 3.0 and Firefox 3.5.
- A new IBM System z Discovery and Auto-Configuration (zDAC) function is added:
 - A new IBM System z Discovery and Auto-Configuration (zDAC), available with z/OS V1.12 on IBM zEnterprise 196 (z196) servers, can help simplify the configuration and reduce the complexity and setup time for new and changed disk and tape I/O configuration changes. zDAC is designed to automatically perform a number of I/O configuration definition tasks for new and changed disk and tape control units connected to a switch or director when attached to a FICON channel. zDAC is designed to discover all IBM disk and tape control units, and other control units that register with the fabric name server. zDAC can save time by discovering new and changed devices for you, and suggesting configurations aligned with best practices for availability and with the I/O policies that you set. For example, with zDAC, adding controllers to an existing 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 that support the function. It proposes new configurations that incorporate the current contents of your I/O definition file (IODF) with additions for new and changed controllers and devices based on a policy you define in the Hardware Configuration Dialog (HCD), which can include preferences for availability and bandwidth, including parallel access volume (PAV), HyperPAV, and DCM specifications, and control unit and device number ranges. These capabilities are integrated with HCD and z/OS Hardware Configuration Manager (HCM) functions.

- The following IBM Health Checker for z/OS enhancements are added:
 - The Health Checker framework allows health checks to be registered without a message table and to issue messages directly without using a message table. This makes it easier to write health checks quickly.
 - The Health Checker framework provides headers to enable you to write health checks using Metal C, in addition to existing support for High-Level Assembler and REXX™. Providing high-level language support can make it easier to write complex health checks. Additionally, sample health checks written in METAL C are provided.
 - SDSF will augment the CK panel by displaying recorded checks on a new health check history panel. The default will be to display up to 10 prior iterations of each check from the log stream, and this support will allow you to browse and print check output from the history panel as you can on the primary CK panel.
 - New health checks are provided for the Parallel Sysplex components, XCF and XES. They are designed to warn you when a coupling facility structure's maximum size as specified in the CFRM policy is more than double its initial size, when any couple data set's (CDS's) maximum system limit is lower than the primary sysplex CDS's system limit, when shared CPs are being used for coupling facility partitions, when the CFRM message-based event management protocol can be used for CFRM event management but the policy-based protocol is being used instead, when your Sysplex Failure Management (SFM) policy does not specify that automatic actions are to be taken to relieve hangs caused by the unresponsiveness of one or more of a CF structure's users, and when a CF does not have a designated percentage of available space to allow for new CF structure allocation, structure expansion, or CF failover recovery. These checks can help you correct and prevent common sysplex management problems.
 - Two new health checks are added for SMB server. The first will detect SMB running in a shared file system environment and alert you that SMB cannot export zFS sysplex-aware read-write file systems in this environment, and the second will determine whether SMB is configured to support the RPC protocol (DCE/DFS) and display a message to remind you that IBM plans to withdraw support for this protocol in a future release.
 - New health checks are added for DFSMS for the communications and active configuration data sets (COMMDS and ACDS) to help ensure these datasets are allocated correctly according to best practices. One new check is designed to alert you that the COMMDS and ACDS are on the same volume. As a best practice, an ACDS and COMMDS must each reside on a volume accessible from all systems in the SMS complex; however, to ease recovery in case of failure, the ACDS should reside on a different volume than the COMMDS. The other new check is intended to identify COMMDS and ACDS data sets that were defined without the REUSE attribute. Specifying REUSE will help to avoid running into space

problems as a result of subsequent ACDS updates or IMPORT/EXPORT functions.

- New health checks are designed for the I/O Supervisor (IOS). IBM recommends using the relatively new MIDAWs and Captured UCB Protection functions introduced in recent releases, and locating eligible I/O-related control blocks above the 16 MB line. These health checks are designed to notify you when these functions are not being used, to help you manage system performance and the use of virtual storage.
- The Health Checker started task supports running with an assigned user ID that has access to the BPX.SUPERUSER profile in the FACILITY class. This will make it unnecessary to run the Health Checker address space with a user ID having UID(0).
- z/OS Communications Server enhances the z/OS Health Checker for z/OS by adding two new checks: one check for IPv4 routing and one check for IPv6 routing. The checks determine whether the total number of indirect routes in the TCP/IP stack routing table exceeds a maximum threshold (default 2000). When this threshold is exceeded, OMPROUTE and the TCP/IP stack can potentially experience high CPU consumption from routing changes. Two new maximum threshold parameters override the default values for the total number of IPv4 and IPv6 indirect routes in a TCP/IP stack routing table before warning messages are issued.
- IBM recommends that you use zFS file systems for z/OS UNIX System Services. In z/OS V1.12, a migration health check identifies HFS file systems you should consider migrating to zFS file systems. This is intended to help you easily obtain and track the list of remaining file systems to be converted. This function is also planned to be available for z/OS V1.11 with the PTF for APAR OA29947.
- SDSF supports displaying information about printers for JES3, and eliminates the requirement for WebSphere® MQ when displaying JES2 MAS-wide data on the initiator panel for JES2 once all systems in the MAS are at z/OS V1.12 JES2. Also, displaying MAS-wide data on the printer panel for JES2 will not require WebSphere MQ when all systems in the JES2 MAS are at or above z/OS V1.11 JES2.
- z/OS V1.12 DFSORT enhancements are designed to:
 - Provide additional diagnostic information intended to help you determine why DFSORT was unable to dynamically allocate all of the requested disk work space and what corrective action may be required.
 - Provide VSAM work space estimates. When VSAM data sets are not closed after updating them, as can happen when a job ends abnormally, the information stored about their sizes is often incorrect. DFSORT is now designed to attempt to estimate the work space needed for VSAM data sets used as input that have not been closed correctly and allocate work data sets accordingly. This is intended to prevent many space-related problems from occurring.
 - Issue diagnostic messages automatically, in error situations, without the need to specify a SORTDIAG DD statement or the DIAGSIM=YES

installation option. This will make it easier to resolve many DFSORT problems without having to re-run a sort solely to collect diagnostic information, and is intended to improve first failure data capture.

- Provide additional information when DFSORT cannot dynamically allocate sufficient work space on a specified volume. New messages are designed to help you resolve these problems quickly.
- Improve reliability when the amount of disk work space required to complete a sort is higher than expected. This can happen when incorrect file size information is provided to DFSORT or when resource contention limits the amount of central storage that can be used as intermediate work space. A new DYNAPCT installation and run-time option allows you to specify additional work data sets to be used if more workspace is required. Also, DFSORT will dynamically adjust secondary extent sizes when work space requirements increase. These changes are expected to reduce sort failures that occur due to insufficient work space.
- DFSMSdfp allows a zFS data set to be recataloged with an extended indirect volume serial using a system symbol. This is designed to allow the zFS file systems used for z/OS system software files (called version root file systems) to be cataloged using extended indirect volume serials, as is possible for non-VSAM data sets, to make cloning and migration easier.
- In z/OS V1.7, Language Environment allowed overridable run-time options to be defined in a new CEEPRMxx member of parmlib. In z/OS V1.12, this is extended to add support for non-overridable (NONOVR) options. With this new support, you no longer need user modifications to change Language Environment run-time options.
- DFSMSrmm provides multiple enhancements in z/OS V1.12. The ACTIVITY file is updated to include the reason why a DFSMSrmm retention limit was reached. This function is also available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA30881. New reports created from the ACTIVITY and extract files are intended to help you see why retention limits were triggered. Also, OPENRULE ignore processing is available for duplicate tape volumes and support is intended to allow you to set a volume hold attribute to prevent expiration and to search and report on volumes which have the hold attribute. It is also intended that the DFSMSrmm ISPF dialog search results can be bypassed when using the CLIST option.
- IDCAMS is enhanced to allow you to delete all members of a partitioned data set in a single operation by specifying a wildcard character (*) as the member name for a data set when using the DELETE command. For example, you can now remove all members of a PDS or PDSE data set in a single command.
- For exploitation in a GDPS® environment with peer-to-peer remote copy (PPRC) controlled DASD devices, HCD and HCM provide a function to automatically generate the OS configuration of the DR site.
- When two or more jobs having the same job name begin to process within the same system clock second and specify the same temporary data set names, the second and subsequent jobs will fail with JCL errors while attempting to allocate data sets with duplicate names. In z/OS V1.12, you will be able to specify, using a

new parmlib option, either that the system is always to use the data set naming convention for unnamed temporary data sets instead, which substantially reduces the probability of this JCL error without the need to change JCL, or that the prior temporary data set naming conventions be used.

- DFSMS provides a system option to control how the system handles multi-volume tape label anomalies. This means that you can now prevent applications from processing tape volumes out of sequence without coding an installation exit.
- The Interactive Storage Management Facility (ISMF), used to manage your SMS configuration, allows you to copy storage group definitions from one control data set (CDS) to another. In z/OS V1.12, you can specify via ISMF that the volume list for pool-type storage groups be copied at the same time. This allows you to copy entire storage groups from one configuration to another without having to add individual volumes to the destination CDS afterward.
- The JESXCF component is changed to allow you to log on to multiple systems within a sysplex using the same TSO/E user ID in a JES2 environment.
- In prior releases, partial release operations for VSAM data sets supported releasing space only on the last volume containing data for each data set. In z/OS V1.12, partial release is extended to support releasing unused volumes in addition to releasing space on the last volume of a multivolume VSAM data set that contains data.
- The IDCAMS DEFINE RECATALOG command is enhanced for multivolume and striped data sets. This new function is designed to automatically create catalog entries with correctly ordered volume lists while eliminating any duplicate volumes that might have been specified. This will make it easier to recatalog multivolume and striped VSAM data sets.
- The z/OS Capacity Provisioning capability can help simplify the provisioning of z/OS resources. For example, it can help automate IBM System z On/Off Capacity on Demand. In z/OS V1.12, the Capacity Provisioning Control Center now supports displaying provisioning reports supported by the Capacity Provisioning Manager. This is intended to simplify the investigation of Capacity Provisioning reports and operation of the Capacity Provisioning server. In addition, the Capacity Provisioning Manager client is updated to provide support for Windows Vista.
- In z/OS V1.9, support was added to write SMF data to log streams. When generating a Postprocessor job using the z/OS V1.12 RMF Postprocessor ISPF interface, users can now specify SMF log streams as an alternative to either SMF data sets or data from the SMF buffers.
- Library Server is designed to improve performance when building the catalog and supporting multiple users on heavily loaded systems. A new Personal Bookcase function in Library Server is intended to allow you to create, use, and share your own subset of the documents from the Library Server catalog. This function is designed to allow you to configure a Personal Bookcase that includes the shelves and documents, as well as the InfoCenters and topics, that you are interested in so you can have the reference documents you routinely use available quickly. Also, indexing will capture the author's intended definition of primary nodes for an InfoCenter's Table of Contents, and administrative improvements include long

filename support, programmatically checking for the required level of Java, and generation of a new Test and Diagnostics page for use by Library Server administrators and IBM support personnel. A number of additional Library Server user interface enhancements are made, including improved navigation between certain dialogs, modernized icons, and descriptive hover popups for documents on a shelf.

- In z/OS V1.12 a new DISPLAY XCF,REALLOCATE,TEST option simulates the reallocation process and provides information about changes the REALLOCATE process that is initiated through the SETXCF START,REALLOCATE command would attempt to make, and any errors that might be encountered if an actual REALLOCATE process were to be performed. This capability will provide information you can use to decide when to invoke the actual REALLOCATE process, and also whether you may need to make any coupling facility configuration changes before issuing the actual REALLOCATE command. A new DISPLAY XCF,REALLOCATE,REPORT command will provide detailed information on the results experienced by a previously executed REALLOCATE command. This capability is intended to help you find such information without searching through the system log for REALLOCATE-related processing and exception messages.
- A number of enhancements are made to the processing of PROGxx parmlib members and to Load Library Lookaside (LLA) processing. These include support in PROGxx for passing a specified parameter to a dynamic exit, automatically including alias names for modules to be placed in Dynamic LPA, and specifying volumes on SYSLIB for data sets so they need not be cataloged in the master catalog; a REPLACE option for exits to assure there is no window during which an exit is unavailable; and a new SVCNUMDEC keyword to specify the SVC number to be added.

Additionally, a new DEFAULTS statement is added, so you can specify processing defaults intended to help prevent common errors. This includes allowing you to specify that LNKLIST DEFINE always require COPYFROM, that it default to COPYFROM(CURRENT), and that it automatically process aliases for modules added to Dynamic LPA.

LLA processing supports the use of dynamic LLA exits and is now able to process multiple MODIFY commands in parallel.

- A new SUMMARY keyword of the DISPLAY SYMBOLS command provides summary information about symbols used on the system, including how many are in use. This can help you determine how many additional symbols can be defined.
- In prior releases, when a corrupted PDSE is detected in the link list during IPL the system often entered a wait state. In z/OS V1.12, the system is designed to attempt to issue a message identifying the corrupt PDSE and continue the IPL without placing that PDSE into the link list. This allows the user to attempt to restore the corrupt PDSE, and re-IPL the system if necessary, without taking a stand-alone dump to debug the problem.

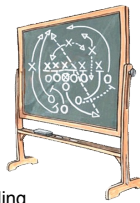
- System Logger corrects the VSAM SHAREOPTIONS for new log stream data sets when it detects that they are not correctly set. Messages will indicate that Logger has detected and corrected a data set's SHAREOPTIONS settings. This new function is intended to prevent data set access problems from arising when SHAREOPTIONS(3,3) has not been set in the data class used to allocate log stream data sets.
- System Logger supports log data set sizes up to 4 GB (from the previous 2 GB limit). This applies to both OFFLOAD and STAGING data set types. As part of this support, System Logger adds messages to show key data set characteristics at allocation and deletion time. This support was also made available for z/OS V1.9, z/OS V1.10, and z/OS V1.11 with the PTFs for APAR OA30548 in February 2010.
- The Web-based CFSIZER tool is designed to help you perform coupling facility structure sizing calculations. Concurrently with z/OS V1.12 availability, CFSIZER provides:
 - More accurate size calculations for the IMS Operations Management (OM) Audit, Resource, and Overflow Sequential Access Method (OSAM) structures
 - Calculated sizes for IBM Session Manager (ISM) Sysplex User structures
 - Improved sizes for XCF signaling structures
 - Calculated sizes for InfoSphere™ Classic control and filter structures
 - Improved sizes for DB2 SCA structures
 - Various usability improvements to the CFSIZER Web pages, including consolidated structure input pages to the OEM cache structure page. For more information about CFSIZER, see
 - <http://www.ibm.com/systems/support/z/cfsizer/>
- A new command is provided in z/OS V1.12 to allow you to specify that all devices on a logical control unit be varied online or offline, or that a path for all devices on a logical control unit be varied online or offline. The new VARY CU command complements the D M=CU command, and makes it easier to change the status of all devices on a control unit, particularly when one is configured in discontinuous device number ranges.

- Slide 9

IBM

Improving Usability and Skills

- **New VARY CU command**
 - New "CU" operand on VARY for "Control Unit" applies to all devices for an LCU
 - No need to find out what range of device numbers apply to the LCU before varying all of them offline or online
 - Before: V (aaaa-bbbb, ... yyyy-zzzz),OFFLINE
 - After: V CU nnnn,OFFLINE
 - (where nnnn = CUNUMBR)
 - Also, support for VARY PATH
- **CFSIZER Improvements**
 - Better size calculations for IMS OM and OSAM, XCF signaling, and DB2 SCA structures
 - Now calculates IBM Session Manager Sysplex User, and InfoSphere™ Classic control and filter, structure sizes
 - Various usability improvements



9
©2010 IBM Corporation


- Slide 10

IBM

Improving Usability and Skills

z/OSMF R12 improvements

- **z/OSMF Incident Log:**
 - Break dumps into multiple data sets that can be sent via FTP in parallel, encrypt incident files
 - Specify additional data sets to send to a vendor
 - Add free-form comments to incidents and FTP destinations in new sortable fields
 - Create of diagnostic log snapshots based on SYSLOG and LOGREC data sets (not just OPERLOG and LOGREC log streams)
- **New z/OSMF interface:**
 - Designed to allow you to add items programmatically to the z/OSMF Navigation tree
- **New WLM policy editor and Monitoring Desktops:**
 - Create, edit, install WLM service definitions, activate WLM policies
 - Monitor WLM status of a sysplex and the systems in a sysplex (requires RMF), and Linux® for System z and Linux on Intel® images (requires the as-is RMF Linux data gatherer tool, rmfpm)
- **z/OSMF Configuration Assistant for z/OS Communications Server now supports configuration for:**
 - Certificate trust chains and certificate revocation lists
 - IKEv2, new cryptographic algorithms and FIPS 140 crypto mode for IPSec and IKE
 - ...and help identify RFC4301 non-compliant IPSec filter rules



z/OSMF Overview
Wednesday 9:30


z/OSMF 1.12 Implementation and Configuration
Wednesday 11:00

Manage Your Workloads and Performance with z/OSMF
Tuesday 4:30

10
©2010 IBM Corporation

Slide 11

Improving Usability and Skills



Multiple TSO/E logons in a JESPLex

- JESXCF used to fail some SSI 80 requests when same ID was logged on multiple systems
- JESXCF changed to make these calls succeed by routing them to the first eligible system
- TSO/E now supports logging on to more than one system in a sysplex

DFSMSrmm™


- Reason retention limit reached added to ACTIVITY file
 - Available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA30881
- New reports created from ACTIVITY and extract files to help you see why retention limits were triggered.
- OPENRULE ignore processing available for duplicate tape volumes
- Support for allowing you to set a volume hold attribute to prevent expiration and to let you search and report on held volumes
- Support to allow you to bypass DFSMSrmm ISPF dialog search results when using the CLIST option

What's New in DFSMSrmm
Wednesday 9:30

11 ©2010 IBM Corporation


Slide 12

Improving Usability and Skills



Library Server Improvements


- New design to improve performance for building catalogs and supporting multiple users on heavily loaded systems
- New Personal BookCase function to allow you to create, use, and share your own subset of documents from a Library Server catalog
 - Configure a Personal BookCase that includes the shelves, documents, infocenters, and topics of interest
- Indexing function to capture the author's intended definition of primary nodes for an InfoCenter's Table of Contents
- Administrative improvements including long file name support, automatic checking for required Java™ levels, and a new Test and Diagnostics page for Library Server administrators
- Revamped user interfaces with improved navigation, modernized icons, and flyover description popups for documents on a shelf



The Future Integrated Information and z/OS Documentation
Tuesday 9:30

12 ©2010 IBM Corporation

Slide 13

Improving Usability and Skills 

New DISPLAY XCF,REALLOCATE,TEST command

- Run "What-If" scenarios for CF structure reallocations
- See what would result from SETXCF START,REALLOCATE
- Also, a new DISPLAY XCF,REALLOCATE,REPORT command shows what the prior REALLOCATE command did

Logger to correct SHAREOPTIONS automatically

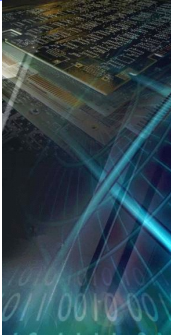
- SHAREOPTIONS(3,3) required
- Will alter new data sets to set required options

RMF improvements

- RMF Postprocessor ISPF interface now allows you to specify SMF log streams, in addition to SMF data sets and SMF buffers
- Postprocessor reports in XML format
- CPU Activity Report shows how many units of work are waiting for a CPU


Health Checker Framework:

- METAL C headers for Health Checker and sample check
- Health check history display in SDSF
- Simpler message interface
- Support for BPX.SUPERUSER (vs. UID(0))




**RMF – The Latest and Greatest
Monday 3:00**

**Exploiting the IBM Health Checker for z/OS Infrastructure
Wednesday 1:30**


13  ©2010 IBM Corporation


Slide 14

Improving Usability and Skills 

New Health Checks:


- XES/XCF checks, to warn you when:
 - A CF structure's maximum size is more than 2X its initial size
 - A CDS's maximum system limit is below the primary sysplex CDS's system limit
 - Shared engines are being used for CF partitions
 - The CFRM message-based event management protocol can be used for CFRM event management but the policy-based protocol is being used instead
 - The SFM policy does not specify that automatic actions are to be taken to relieve hangs caused by the unresponsiveness of one or more of a CF structure's users
 - A CF does not have a designated percentage of available space to allow for new CF structure allocation, structure expansion, or CF failover recovery.
- SMB checks:
 - Detect shared file system environment and alert you that zFS sysplex-aware read-write file systems cannot be exported
 - Warn you when SMB is configured to support the RPC protocol (DCE/DFS)



14  ©2010 IBM Corporation


Slide 15

Improving Usability and Skills



More New Health Checks:


- **DFSMS™**
 - Warn you if COMMDS and ACDS data sets are on the same volume
 - Identify COMMDS and ACDS data sets defined without the REUSE
- **IPv4 and IPv6 indirect routes**
 - Warn when the threshold is exceeded, which can cause high CPU consumption
- **HFS in use**
 - IBM recommends migrating to zFS
- **IOS**
 - Warn you if MIDAW and Captured UCB Protection functions are not being used, and when I/O-related control blocks eligible to be placed above the 16 MB line are below the line
- **Communications Server**
 - Too many indirect routes for IPv4 or IPv6 (default threshold is 2,000)



15 ©2010 IBM Corporation


Slide 16

Improving Usability and Skills




Dynamic LPA and Link List, LLA exit, and Dynamic Exit improvements

- Automatic alias processing option in PROGxx and on CSVDYLPA (default will automatically add aliases)
- VOLUME specification for SYSLIB in PROGxx
- Dynamic Exit support for passing a 1-8 byte parameter to an exit via PROGxx, SETPROG, and CSVDYNX
- New PROGxx DEFAULTS statements
 - Whether COPYFROM is required for link list
 - Whether to default to COPYFROM(CURRENT) for link list
 - Whether to add LPA aliases automatically
- LLA exits to support dynamic exit facility
- MODIFY LLA to accept multiple commands
- New SVCNUMDEC option for LPA ADD
- Dynamic Exit support for REPLACE



16 ©2010 IBM Corporation

Slide 17

Improving Usability and Skills 

CEEPRMxx and SETCEE support for NONOVR

- CEEPRMxx was introduced in z/OS R7 to supplant CEEExOPT, etc., usermods
- Overridable (OVR) options were originally supported
- Support added for non-overridable (NONOVR) options in z/OS R12

What's New in Language Environment for z/OS Monday 3:00

SDSF enhancements

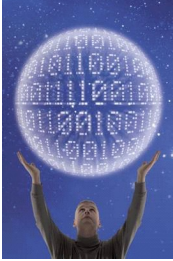
- Printer displays for JES3
- MQ requirement removed for JES2 initiator displays
 - All systems in the JES2 MAS must be at z/OS V1.12 JES2
- MQ requirement removed for JES2 device displays
 - All systems in the MAS at or above z/OS V1.11 JES2


“Tool and Toy” tsocmd function added to z/OS UNIX

- Support for stdin and stdout
- New return codes for failing TSO/E commands
- TSOPROFILE environment variable


New D SYMBOLS SUMMARY command

- Displays information about the number of symbols in use
- D SYMBOLS DETAIL the same as D SYMBOLS without SUMMARY



17  ©2010 IBM Corporation

Slide 18


Improving Usability and Skills 

DFSMS Improvements


- **PDS and PDSE “delete all members” support**
 - DELETE DATA.SET.NAME(*)
 - e.g., for IDCAMS DELETE GDG FORCE
- **Support for indirectly cataloged zFS data sets:**
 - DELETE NOSCRATCH
 - DEFINE CLUSTER (NAME(name) LINEAR VOLUMES(&symbol) RECATALOG)
 - Does not support other VSAM (including linear) data sets, only zFS
- **DEFINE RECATALOG**
 - Build catalog entries for multivolume data sets in the correct order automatically
- **Partial Release (PARTREL) improvement for VSAM**
 - Support for releasing unused volumes
 - In addition to releasing space on the last volume of a multivolume VSAM data set that contains data
- **ISMF storage group copy improvements**
 - ISMF allows storage group definitions to be copied from one CDS to another
 - Extended to let you specify that volume lists for pool-type storage groups be copied
 - Designed to allow you to copy entire storage groups from one configuration to another without having to add their volumes to the destination CDS afterward

MVS™ Storage Management Project Opening Monday 1:30

What's new in DFSMS Monday 3:00

18  ©2010 IBM Corporation

Slide 19

Improving Usability and Skills 


DFSORT Improvements

MVS Storage Management Project Opening Monday 1:30

- Better diagnostic information when DFSORT was unable to dynamically allocate requested disk work space
- DFSORT now designed to try to estimate the work space needed for VSAM data sets used as input that were not been closed correctly
- Diagnostic messages without a SORTDIAG DD statement (or the DIAGSIM=YES installation option) for many errors
- New recovery option for when work space required to complete a sort is higher than expected
 - New DYNAPCT installation and run-time options
- Dynamically adjust secondary extent sizes when work space requirements increase
- Allow you to specify that DFSORT use memory object storage instead of Hiperspace™ for intermediate work storage using a new MOWRK option


19 ©2010 IBM Corporation

Slide 20

Improving Usability and Skills 

Help With Duplicate Temporary Data Set Names

- Everyone's favorite messages:
 - IGD17001I DUPLICATE DATA SET NAME ON VOLUME volser FOR DATA SET data.set.name
 - IEF453I jobname - JOB FAILED - JCL ERROR - TIME=hh.mm.ss
- Temporary data set name format if you code &&tempname:
 - SYSyyddd.Thhmmss.RA000.jobname.dsetname.Hgg
- Temporary data set name if you do not code a name:
 - SYSyyddd.Thhmmss.RA000.jobname.Rggnnnnn
- When multiple jobs with the same name start within the same clock second, IGD17001I can occur when they both try to allocate a data set:
 - With the same name
 - On the same volume
- New ALLOCxx (and SETALLOC) setting:
 - SYSTEM TEMPDSFORMAT(UNIQUE|INCLUDELABEL)
- UNIQUE specifies that the system should always use the second format
- Vastly decreases the chances of a naming collision causing an allocation failure because Rggnnnnn is changed for each allocation



20 ©2010 IBM Corporation

Scalability and Performance

EAV Support for 'most everything else'

- z/OS R10 introduced EAV with support for VSAM (incl. zFS)
- z/OS R11 added Extended Format Sequential and support for data sets spanning the 64K cylinder line
- z/OS R12 adds:
 - PDS and PDSE (including load modules and program objects)
 - Plain vanilla (nonextended format) sequential
 - BDAM
 - GDG
 - LPALIB, LPA list, link list data sets, SYSn.IPLPARM, SVCLIB
 - Catalogs, VVDSs
 - JES2 and JES3 spool and checkpoint, JES3 JCT
 - DFSMSmm, DFSMSHsm™ data sets
 - Standalone Dump data set and AMASPZAP support
 - VSAM AIX support in Language Environment
- **No** support for above the line for:
 - Imbed and Keyrange attributes, incompatible CA sizes for VSAM
 - NUCLEUS, SVCLIB, LOGREC, VTOC, VTOCIX,
 - RACF databases, Page data sets, HFS data sets
 - Parmlib concatenation data sets
 - XRC Control, Master or Cluster non-VSAM data sets

21
©2010 IBM Corporation

Scalability and performance

IBM System z has long understood the balance between scalability and performance and efficiency of the platform. The major components of the system, the processors, storage, I/O, and software, work together and help manage system resources. Essentially, z/OS and its subsystems provide for scalability not only based on faster chip speeds, but also via efficient single-image n-way processor growth, highly scalable sysplex clustering for horizontal growth, and scalable storage and data management as well.

z/OS V1.12 in and of itself has had many performance improvements for your key workloads:

CICS VSAM workload performance could improve by up to 44% with the use of VSAM Control Area (CA) Reclaim. The actual benefit may be more or less depending on the degree of fragmentation of VSAM KSDS. It is anticipated that rarely reorganized, severely fragmented data sets will see more benefit.

DB2 9 for z/OS startup time can be reduced significantly. The reduction in DB2 9 for z/OS startup time is through z/OS Allocation, DFSMSdfp, and GRS improvements; actual benefit will be a large range and will depend on number of data sets opened. It is anticipated that address spaces opening hundreds of thousands of data sets will see more benefit.

SVC dump capture time can be reduced by up to 50-90%. Actual SVC dump time will depend on amount of data being captured and the amount of that data dumped from auxiliary storage.

z/OS XML System Services parsing performance is improved. For example, validation parsing performance may be improved by 30%-50%. Actual benefit will depend on the amount of data being parsed and the degree of complexity of the schema. In addition, a revolutionary XML validation performance feature - new fragment validation - enables you to validate only a portion of a document. By revalidating only the fragments being updated, DB2 10 for z/OS pureXML can avoid the costly revalidation of entire XML documents, which without this function can take many times longer.

Network throughput is improved. Network CPU overhead for Application Transparent - Transport Layer Security (AT-TLS) can be improved by 30%. The AT-TLS CPU consumption results were obtained on System z10, model 2097-E64. Actual AT-TLS CPU consumption improvement will depend on amount of data being transmitted and whether the workload is interactive or streaming.

Network throughput for interactive workloads can be improved by 30-50% using the new Inbound Workload Queueing (IWQ) function, which is exclusive to OSA-Express3 on z196 and System z10 servers. The interactive networking throughput measurements were obtained on System z10, model 2097-E64. Actual benefit will depend 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.

z/OS V1.12 and IBM System Storage® provide vast improvements in scalability and data-handling capabilities:

- With z/OS V1.12, DFSMS supports most additional data set types and data sets in Extended Address Volumes (EAVs). EAVs can help relieve storage constraints as well as simplify storage management by providing the ability to manage fewer, large volumes as opposed to many small volumes.
- The IBM zEnterprise System adds additional scalability and performance capabilities as well.
 - 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 zEnterprise 196 Model M80 provides up to 1.6 times the total system capacity and 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 can provide a significant performance improvement for Java workloads.
 - DB2 9 for z/OS is the first exploiter of the IBM zEnterprise BladeCenter optimizers. The IBM Smart Analytics Optimizer is designed to help mitigate database administration costs while speeding up the performance of typical data warehouse or OLAP-specific workloads by an order-of-

magnitude. Actual query performance improvement will depend on the specific query workload.

- z/OS V1.12 supports three subchannel sets on IBM zEnterprise 196 servers. This helps relieve subchannel constraints, and can allow you to define larger I/O configurations that include large numbers of Metro Mirror (PPRC) secondaries and Parallel Access Volume (PAV) aliases.

Details on the performance and scalability enhancements for z/OS V1.12:

VSAM performance improvement

Over time, VSAM key-sequenced data sets (KSDS) for which records are added and deleted have often become fragmented and have a significant number of empty Control Areas (CAs) that consume DASD space, increase the size of the indexes, and reduce performance. In z/OS V1.12, DFSMSdfp allows you to specify that VSAM dynamically reclaim empty control areas for KSDSs, including catalogs and those used for record-level sharing (RLS), and reclaim the associated index records as needed. This new function is intended to help you preserve performance, minimize space utilization for KSDSs, and improve application availability, and allow you to avoid the need to reorganize most KSDSs. Depending on how fragmented the data set is, and how the data set is defined and accessed, you may get a very wide range of performance results. VSAM performance improvement may be up to 44% or more for rarely reorganized, severely fragmented data sets and will depend on the degree of VSAM key-sequenced data set (KSDS) fragmentation your applications encounter without this function. It is anticipated that customers who do not periodically reorganize their severely fragmented VSAM data will see the most benefit.

Reduced startup time

Design changes for Allocation, DFSMSdfp, and GRS are intended to improve performance for address spaces that allocate a large number of data sets in a short time. These changes are expected to markedly reduce the startup time for these address spaces, such as DB2 address spaces and batch jobs that process a large number of data sets per job step. For example, the reduction in DB2 9 for z/OS startup time will depend on number of data sets opened. It is anticipated that address spaces opening tens of thousands of data sets will see more benefit.

Improved SVC dump capture time

In z/OS V1.12, RSM and dump processing are designed to improve capture performance for large amounts of data during SVC dump processing. This is expected to markedly reduce dump capture time when a large amount of data must be paged in during SVC dump processing. Internal IBM laboratory tests have shown that the capture time for SVC dumps can be reduced by up to 50-90% for large dumps with a substantial percentage of data on auxiliary storage. The amount of improvement is expected to vary depending on

how much data must be paged in during dump processing, how much real storage is available to the system, and the system's workload.

XML validation performance improvements

z/OS XML System Services parsing performance is improved. For example, validation parsing performance may be improved by 30%-50%. Actual benefit will depend on the amount of data being parsed and the degree of complexity of the schema. In addition, a revolutionary XML validation performance feature - new fragment validation - enables you to validate only a portion of a document. By revalidating only the fragments updated, DB2 10 for z/OS pureXML can avoid the costly revalidation of entire XML documents, which without this function can take many times longer.

Additional z/OS V1.12 scalability and performance improvements:

- DFSMS now supports additional data set types, including basic and large format sequential data sets, partitioned (PDS/PDSE) data sets, direct (BDAM) data sets and catalogs in the extended addressing space (EAS) on extended address volumes (EAVs). Support is also included for generation data groups (GDGs) and VSAM volume data sets (VVDSs). Overall, EAV helps you relieve storage constraints as well as simplify storage management by providing the ability to manage fewer, large volumes as opposed to many small volumes. Other new EAV-related functions include:
- Support is added to make all data sets used by DFSMSrmm eligible for allocation in the extended addressing space of an EAV. This includes the DFSMSrmm journal and dynamically allocated temporary files.
- DFSMSHsm now supports additional non-VSAM data set types in the EAS on EAVs. All EAV volume space can be used on DFSMSHsm owned disk volumes, including ML1, ML2, and backup volumes. Migration copies and backup versions can be allocated in the EAS. Additionally, DFSMSHsm owned data sets, including the journal, logs, CDS backup copies, PDA, and fast replication catalog information data sets, can now reside in the EAS on EAVs.
- Language Environment provides support for C/C++ to access alternate indexes (AIXs) for extended format VSAM key-sequenced data sets (KSDS) that reside in the EAS on an EAV.
- Support is provided for additional input and output data set types including sequential (both basic and large format) and partitioned (PDS/PDSE) data sets in the EAS on EAVs.
- Support is provided for additional work data set types of basic and large format sequential in the EAS on EAVs. This allows you to specify sort work data sets up to 1,048,574 tracks in size to support larger sort operations.
- Superzap (AMASPZAP) supports dumping and altering data for sequential, partitioned, and direct data sets placed in EAS on EAVs.
- JES2 is designed to allow both spool and checkpoint data sets to reside in the EAS on an EAV, making it possible to place both spool and checkpoint data sets


- anywhere on an EAV and to define spool data sets up to the maximum size of 1,000,000 tracks (approximately 56 GB).
- JES3 is designed to allow spool, checkpoint, and Job Control Table (JCT) data sets to be placed anywhere on an EAV.
 - Stand-alone Dump is designed to support extended format dump data sets in the EAS on EAVs.
 - z/OS V1.12 supports three subchannel sets on IBM zEnterprise 196 (z196) servers. This helps relieve subchannel constraints, and can allow you to define larger I/O configurations that include large numbers of Metro Mirror (PPRC) secondaries and Parallel Access Volume (PAV) aliases. As with the prior support for two subchannel sets, you can define base devices, aliases, and secondaries in the first subchannel set (set zero), and define only aliases and secondaries in subchannel sets one and two. All three subchannel sets support ESCON®, FICON, and zHPF protocols. This support is also available on z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA30677.
 - z/OS V1.12 on zEnterprise 196 servers with Coupling Facility Control Code (CFCC) Level 17 supports up to 2047 structures per Coupling Facility (CF) image, up from the prior limit of 1023. This allows you to define a larger number of data sharing groups, which can help when a large number of structures must be defined, such as to support SAP configurations or to enable large Parallel Sysplex configurations to be merged. This function requires the PTF for APAR OA32807; PTFs are also available for z/OS v1.10 and z/OS V1.11.
 - z/OS V1.12 on zEnterprise 196 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. This support requires the PTF for APAR OA32807; PTFs are also available for z/OS V1.10 and z/OS V1.11.
 - z/OS V1.12 supports larger Coupling Facility (CF) structures. The maximum size you can specify for a CF structure is increased from slightly below 100 GB (99,999,999 KB) to 1 TB. Also, the CFRM policy utility (IXCMIAPU) is updated to allow you to specify structure sizes in units of KB, MB, GB, and TB. These changes improve both Parallel Sysplex CF structure scalability and ease of use.
 - In z/OS V1.12, DFSORT can now use memory object storage instead of Hiperspace™ as intermediate work storage. New installation default and run time option MOWRK can be used to control whether DFSORT can use memory objects as intermediate work storage. New DFSORT messages indicate whether memory object storage was used as main storage or work storage.
 - In z/OS V1.12 a program that invokes DFSORT, ICETOOL, or ICEGENER can dynamically allocate input, output, and work data sets using the options for XTIOs, uncaptured UCBs, and DSABs above 16 MB.
 - Some workloads require an increasing number of open data sets. In z/OS V1.12, the BSAM, QSAM, and BPAM (basic and queued sequential, and basic partitioned access methods), OPEN/CLOSE/EOV, common VTOC access facility (CVAF), direct access device storage manager (DADSM), EXCP (execute channel program), and RACF processing support the use of an extended task I/O

- table (XTIOT) with uncaptured UCBs, and support data set association blocks (DSABs) above the 16 MB line. This is expected to allow more data sets to be allocated by an address space and to provide virtual storage constraint relief for programs using large numbers of DASD and tape data sets. Also, these functions have been updated to provide support:
- SNAP/SNAPX services and dump processing (including that for SVC, SYSABEND, SYMDUMP, and SYSUDUMP)
 - AMASPZAP
 - The Program Management Binder
 - TSO/E
 - DFSORT
- OAM provides API support for the Object Storage and Retrieval function (OSR) to run in a CICS threadsafe environment. This is intended to allow exploiters to take advantage of the improved multitasking and throughput capabilities provided by threadsafe programming. Additionally, the Volume Recovery utility is designed to improve performance in certain situations when recovering object data stored on optical and tape media. Improvements are expected to be most noticeable when recovering a backup volume containing objects with primary copies in a large number of different collections on a large number of different volumes.
 - Large (1 MB) pages were introduced in z/OS V1.10 on IBM System z10 servers. In z/OS V1.12, the system is designed to back part of the nucleus data area using 1 MB pages, when they are available. This is intended to reduce the overhead of memory management for nucleus pages and to free translation lookaside buffer (TLB) entries so they can be used for other storage areas. This is expected to help reduce the number of address translations that need to be performed by the system and help improve overall system performance.
 - In z/OS V1.7, support was introduced in DFSMSdftp for large format sequential data sets (DSNTYPE=LARGE). In z/OS V1.8, Language Environment added support for these data sets using noseek (QSAM). Support for seek (BSAM) was limited to data sets no larger than 64K tracks on any volume when opened for read. In z/OS V1.12, seek (BSAM) support is extended to data sets up to the maximum size when using record I/O. Binary and text I/O with seek continue to be supported for data sets up to 64K tracks in size on any volume when opened for read.
 - DFSMS now supports catalogs with extended addressability (EA). This is designed to make it possible to define and use Integrated Catalog Facility (ICF) catalogs larger than 4 GB, up to the current maximum size of a volume. The new architectural limit on catalog size is approximately 140 TB.
 - z/OS Communications Server AT-TLS processing is designed to provide reduced CPU usage for encryption and decryption of application data while improving throughput for some types of workloads. This function is automatically enabled.
 - In z/OS V1.12 Communications Server, processing for OSA-Express in QDIO mode supports inbound workload queueing. Inbound workload queueing uses multiple input queues for each QDIO data device to improve TCP/IP stack scalability and general network optimization.

- When inbound workload queueing 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 queueing 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 Communications Server introduces sysplex distribution of IPv6 connections to non-z/OS targets, similar to the sysplex distribution of IPv4 connections to non-z/OS targets introduced in z/OS V1R11 Communications Server. An IBM WebSphere DataPower® appliance is currently the only non-z/OS target that supports sysplex distributor load balancing. DataPower appliances are often used as a front-end processing tier for z/OS applications, allowing for more efficient handling of Web services for a second tier of z/OS applications. When the DataPower tier completes the handling of a request, DataPower can route the request to a tier 2 sysplex distributor, which can load balance the request to the second tier of z/OS applications.
- z/OS Communications Server TN3270E Telnet server provides access method control block (ACB) sharing for logical units (LUs) as a way to help reduce ECSA usage. Prior to z/OS V1.12 Communications Server every Telnet LU name opened its own ACB to VTAM. You can code a new SHAREACB statement to allow multiple Telnet LUs to share a single ACB, which reduces the overall amount of ECSA (and Telnet private) storage allocated to support Telnet sessions.
- z/OS V1.12 Communications Server extends the SNMP manager API to support these tasks:
 - Create and retrieve SNMP values of type UNSIGNED32.
 - Configure an authoritative engine ID for SNMPv3 traps. Currently, the SNMP manager API creates its own SNMPv3 authoritative engine ID, part of which is a randomized value. A configured authoritative engine ID can be used with SNMP trap receiver applications so that the trap receiver applications recognize specific SNMP manager API applications, when processing SNMPv3 traps.
- VSAM record level sharing (RLS) supports striped data sets. This brings the benefits of VSAM striping, such as allowing single application requests for records in multiple tracks or control intervals (CIs), to be satisfied by concurrent I/O requests to multiple volumes.
- DFSMSdss is designed to use larger blocks when possible for DUMP, COPYDUMP, and RESTORE operations, and to support Extended Format Sequential dump data sets on DASD for DUMP, RESTORE, and COPYDUMP. The use of larger block sizes is intended to improve performance for these operations, and using Extended Format dump data sets is intended to support striping and compression.
- DFSMShsm is designed to support parallel processing for recovery from dump tape volumes when the dumps reside on multiple tape volumes and multiple tape drives are available. This new function is intended to allow you to specify that up


- to 64 concurrent tasks be used to help speed recovery processing. Also, this is designed to allow you to restore copy pools from tape using DFSMSHsm.
- DFSMSHsm Space Management is designed for performance improvements. Primary Space Management, Interval Migration, and the MIGRATE PRIMARY command will now overlap volume preprocessing with volume data movement in order to reduce elapsed time.
 - The DFSMSHsm DUMP function used to copy source disk volumes to a target tape volume is enhanced. The dump stacking function is designed to allow up to 255 source volumes to be dumped to a single tape volume, up from the prior limit of 99. This is intended to help you take better advantage of large capacity tape cartridges.
 - The Catalog address space (CAS) is designed to check for SYSZTIOT enqueue contention periodically. Based on an interval you specify and the reason for contention, CAS is designed to write a logrec record and a notification message when tasks have waited longer than the specified interval and contention checking is active. A new MODIFY CATALOG,CONTENTION command allows you to specify a different interval than the 10-minute default or to disable CAS contention detection. This new function is intended to warn about tasks that take excessive time to complete, or never complete, from affecting Catalog performance.
 - Language Environment is designed to improve performance for string manipulation intensive applications, such as certain applications written in the Perl language.
 - The Infoprint Server removes several constraints in z/OS V1.12. IP PrintWay™ Extended Mode now supports as many SYSOUT data sets per job as are supported by the system, and as many jobs as are supported by the JES2 or JES3. IP PrintWay Extended Mode now supports the Response Notification exit. Also, the Infoprint Server Line Printer Daemon (LPD) now supports larger file sizes over 2 GB; this support is also available on z/OS V1.10 and higher with the PTF for APAR OA28795. Last, Infoprint Server is now designed to give higher priority to despooling and printing existing jobs than receiving new work from JES2 or JES3. This is intended to help relieve constraints and reduce spool occupancy of active Infoprint Server jobs.
 - Two new services based on existing XCF signaling services are introduced to support the use of 64-bit addressable virtual storage message buffers and associated input and output parameters. The two new services, IXCMSGOX and IXCMSGIX, are the 64-bit counterparts of the existing IXCMSGO and IXCMSGI services. These new services make it easier for exploiters to achieve virtual storage constraint relief by removing the need to copy message buffers and associated storage structures from 64-bit addressable virtual storage to 31-bit storage and back.

Slide 22

Scalability and Performance 


- **XL C/C++ Enhancements for z196 Servers**
 - Benefits from Out-of-Order (OOO) Execution
 - New ARCH(9) and TUNE(9) support
 - Support for new instructions from the load/store-on-condition and the distinct-operands facilities
 - 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.

What's New in Enterprise PL/I V3R9 and C/C++ V1.12
Monday 4:30




22 ©2010 IBM Corporation

Slide 23

Scalability and Performance* 


- **Support for three subchannel sets (0, 1, and 2)**
 - On z196 servers
 - Aliases can be defined in any subchannel set
 - PPRC secondaries must all be in subchannel set 1 or in subchannel set 2.
 - Aliases can be defined in the same subchannel set as the PPRC-pair secondary devices
- **Parallel subsystem initialization**
 - All subsystems named after a new BEGINPARALLEL keyword in IEFSSNxx will start in parallel
 - Intended to reduce restart time
- **BCP exploitation of Large Pages for NUCLEUS**
 - Reduced TLB residency for NUCLEUS pages makes TLB entries available for other things
- **DFSMSdss™ Dump/Restore/Copydump performance**
 - Use 256K blocks rather than 64K blocks
 - Expected to improve throughput for some operations

What's New in DFSMSdss and System Data Mover
Wednesday 11:00




23 ©2010 IBM Corporation


Slide 24

Scalability and Performance 


- **Larger Catalogs**
 - DFSMS support for catalogs with extended addressability
 - Designed to support ICF catalogs larger than 4 GB
 - New maximum is same as the volume size, currently 223 GB
- **Parallel processing for DFSMSHsm**
 - Parallel recovery from multiple dump tapes
 - Specify that up to 64 concurrent tasks be used
 - Also designed to allow you to restore Fast Recovery copy pools from tape using DFSMSHsm
 - Parallel processing for Primary Space Management, Interval Migration, and Command Volume Migration
 - Specify using a new option
- **DFSMSHsm tape stacking**
 - Dump stacking maximum increased from 99 dumps per tape to 255
 - Intended to improve utilization for large-capacity tape media

What's New in DFSMSHsm
Tuesday 11:00



24 


Slide 25

Scalability and Performance 

Parallel Sysplex Scalability


- **Support for up to 2047 structures per CF image**
 - Up from prior limit of 1023
 - Allows you to define more sharing groups
 - Support large SAP configurations
 - Enable large Parallel Sysplex configurations to be merged
 - Requires the PTF for APAR OA32807 and a z196 server with CFCC Level 17
 - PTFs are also available for z/OS R10 and z/OS R11
- **Larger CF structures**
 - New maximum size is 1 TB
 - Old maximum was 99,999,999 KB
 - IXCM/APU updated to support unit specifications in KB, MB, GB, and TB
- **Support for more connectors to list and lock structures**
 - XES and CFCC already support 255 connectors to cache structures
 - XES now supports up to:
 - 247 connectors to a lock structure
 - 27 connectors to a serialized list structure
 - 255 connectors to an unserialized list structure
 - Requires the PTF for APAR OA32807 and a z196 server with CFCC Level 17
 - PTFs also available for z/OS R10 and z/OS R11

z/OS Parallel Sysplex: Update
Tuesday 11:00


25 

Slide 26

Scalability and Performance




- **VSAM/RLS Striping**
 - Extend striping for VSAM to RLS
 - Usual advantages of striping available for RLS
- **XTIOT support in:**
 - BSAM, BPAM, QSAM, EXCP, DADSM, CVAF
 - SVC dump, SYSABEND, SYSMDUMP, and SYSUDUMP
 - AMASPZAP, Binder, TSO/E, RACF, DFSORT
 - Support for uncaptured UCBs, DSABs above the line
 - Provides 24-bit VSCR for exploiters
- **Bigger offload and staging data sets for System Logger**
 - Raise limit from 2 G to 4 G
 - Also, new messages to show key data set characteristics at allocation and deletion time
 - Available for z/OS V1.9 and higher with PTFs for APAR OA30548



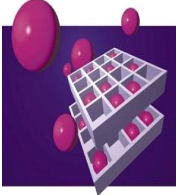
26 ©2010 IBM Corporation

Slide 27


Scalability and Performance



- **OAM**
 - API support for the Object Storage and Retrieval function (OSR) to run in a CICS threadsafe environment
 - Volume Recovery utility will be designed to improve performance in certain situations when recovering object data stored on optical and tape media.
 - Expected to be most noticeable when recovering a backup volume containing objects with primary copies in a large number of different collections on a large number of different volumes
- **Language Environment**
 - BSAM (seek) support for large format sequential data sets (DSNTYPE=LARGE) when using record I/O
 - Binary and text I/O with seek continue to be supported for data sets up to 64K tracks in size on any volume when opened for read



27 ©2010 IBM Corporation



Improving Availability


VSAM Control Area (CA) Reclaim

- For VSAM key-sequenced data sets (KSDS) including VSAM RLS data sets *and catalogs*
- Designed to reuse empty CAs and empty index records in place rather than extending the data set
- Intended to eliminate need to copy/delete/reload VSAM data sets to reclaim space taken up by empty CAs, rebuild index, and restore performance
- Substantial performance improvements possible*
- Toleration on z/OS R10 and z/OS R11 via APARs OA25108, OA26256, OA26466, and OA27557
 - *Be sure the PTFs are on before turning on CA Reclaim!*

Using CA Reclaim:

- Parameter in IGDSMSxx: CA_RECLAIM(NONE | DATACLAS)
- SETSMS support
- DATACLAS attribute: CA_Reclaim(Y/N)
- IDCAMS ALTER support for individual data sets
- LISTCAT shows whether CA Reclaim is being used for a data set
 - With the PTF for APAR OA33315, it will also show how many CAs have been reclaimed
- SMF64 records include a Reclaim count field, SMF64DAU

z/OS 1.12 New Feature VSAM CA Reclaim
Monday 4:30



28 *IBM system tests indicate performance of dedicated VSAM workload could improve by up to 44%. Based on IBM Lab results, your results will vary. Actual benefit may be more or less and will depend on the degree of VSAM data fragmentation and how the data is accessed. It is anticipated that VSAM key-sequenced data sets (KSDS) that are heavily fragmented for evenly indexed will see the most benefit. For applications that delete a large number of records from a narrow key range and then immediately re-insert them, CA Reclaim could result in some performance degradation. ©2010 IBM Corporation

Availability

Resilience that helps reduce risk from outages

There is more to "availability" than just the server being up -- the applications and the data must be available as well. For the System z platform this means hardware, I/O connectivity, operating system, subsystem, database, network, and application availability, too. The System z hardware is designed to reduce planned and unplanned outages through the use of self-healing capabilities, redundant componentry, dynamic sparing, and the ability for concurrent upgrades and microcode changes. Data availability and integrity are upheld with capabilities such as address space isolation, storage protect keys, I/O channel redundancy, and I/O error checking.

Beyond the single system is z/OS Parallel Sysplex clustering (see also the Scalability and performance section). Parallel Sysplex clustering is designed to provide your data sharing applications and data with not only continuous availability for both planned and unplanned outages, but also near-linear scalability and read/write access to shared data across all systems in the sysplex for data sharing applications.

z/OS V1.12 has numerous updates that improve the availability of specific applications and workloads. VSAM Control Area (CA) Reclaim is intended to help you preserve performance, minimize space utilization for many KSDSs, and improve application availability, and allow you to avoid many or most application outages needed to reorganize KSDSs. RSM and dump processing are designed to improve capture performance for large amounts of data during SVC dump processing, reducing the potential for system disruptions. Allocation, GRS, and DFSMSdfp updates are designed to improve startup performance for address spaces that allocate a large number of data

sets in a short time, helping to shorten restart time for DB2 and batch jobs that process a large number of data sets.

z/OS also has error checking, fault tolerance, isolation, error recovery, and Parallel Sysplex capabilities that it continues to enhance every year. Unlike other operating systems, z/OS makes advances in a new, innovative direction for availability. z/OS extends its proactive learning, monitoring, and analysis, to enable the system to analyze a component or a subsystem that it suspects is failing and provide warnings and guidance for operators and system programmers.

z/OS extends its high availability characteristics by going beyond failure detection to predicting potential problems before they occur. With Predictive Failure Analysis (PFA), introduced with z/OS V1.11, your z/OS system is designed to learn heuristically from its own environment to anticipate and report on system abnormalities. It compares present and past behaviors and models system behavior in the future, and is intended to notify you when a system trend might cause a problem. The initial delivery of this function is also available on z/OS V1.10 with PTFs UA46241 and UA46243. For z/OS V1.12

Predictive Failure Analysis (PFA) monitors the rate at which the system is generating SMF records. When the rate is abnormally high, the system is designed to issue an alert warning you of a possible problem, potentially avoiding an outage. PFA can take into account the normal swings of daily, weekly, or monthly spikes and can learn the idiosyncrasies of your system, thus avoiding false warnings given by static monitors. In z/OS V1.12, a new capability, z/OS Runtime Diagnostics, is designed to help when the need for quick decision-making is required. With Runtime Diagnostics, your z/OS system is designed to analyze key system indicators of a running system. The goal is to help you identify the root of problems that cause system degradation on systems that are still responsive to operator commands. Runtime Diagnostics is anticipated to run quickly to return results fast enough to aid you in making decisions about alternative corrective actions and facilitate high levels of system and application availability.

IBM Tivoli OMEGAMON® XE on z/OS V4.20 (5698-A33), IBM Tivoli NetView® for z/OS V5.4 (5697-ENV), and IBM Tivoli System Automation for z/OS (5698-SA3) are updated to provide performance, availability monitoring, and automation of complex and repetitive tasks. OMEGAMON XE on z/OS is integrated with z/OS RMF and Health Checker, and with z/OSMF for performance and operations management. NetView for z/OS and System Automation for z/OS can automate responses to the corrective actions identified by Predictive Failure Analysis and Runtime Diagnostics to help increase the reliability, availability, and serviceability of applications and subsystems on z/OS. In z/OS V1.12, a new Timed Auto Reply function enables the system to respond automatically to write to operator with reply (WTOR) messages. This new function is designed to help provide a timely response to WTORs and help prevent delayed responses from causing system problems.

z/OS availability is beyond the server as well. Parallel Sysplex can provide a large single system image, dynamic load balancing, fault tolerance, and automatic restart capabilities,

so a single cluster can be used for scalability and performance as well as for availability and disaster recovery. With z/OS V1.12, Parallel Sysplex technology is planned to be updated with new health checks; improved command routing; and improved network traffic routing, security, availability, and reporting. There are also plans to provide autonomies whereby the z/OS system can help identify CF structures and Sysplex Distributor network connections that are unresponsive or in a degraded state. In addition, GRS and XCF components are designed to automatically initiate actions to preserve sysplex availability to help reduce the incidence of sysplex-wide problems that can result from unresponsive critical components.

Details on availability improvements for z/OS V1.12 include:

- Improved VSAM availability. Over time, VSAM key-sequenced data sets (KSDS) for which records are added and deleted have often become fragmented and have a significant number of empty Control Areas (CAs) that consume DASD space, increase the size of the indexes, and reduce performance. Performance and DASD space utilization can usually be improved for such data sets by copying, deleting and reallocating, and reloading them. This requires scheduled outages for applications using these data sets. In z/OS V1.12, DFSMSdftp allows you to specify that VSAM dynamically reclaim empty control areas for KSDSs, including those used for record-level sharing (RLS), and reclaim the associated index records as needed. This new function is intended to help you preserve performance, minimize space utilization for KSDSs, and improve application availability, and allow you to minimize the need to reorganize KSDSs.
- In z/OS V1.12, a new component named z/OS Runtime Diagnostics helps you reduce the time spent deciding what actions to take to resolve a problem. It can be used to identify potentially related symptoms and causes when it appears a significant system problem might affect the system's ability to process your workloads. Often, you must quickly analyze these problems to preserve application availability. Runtime Diagnostics runs using the START operator command and returns results quickly to help you decide among alternative corrective actions and maintain high levels of system and application availability.
- A new Timed Auto Reply Function provides an additional way for the system to respond automatically to write to operator with reply (WTOR) messages. This new function allows you to specify message IDs, timeout values, and default responses in an auto-reply policy, and to be able to change, activate, and deactivate auto-reply with operator commands. Also, when enabled, it starts very early in the IPL process, before conventional message-based automation is available, and continues unless deactivated. You can also replace or modify an IBM-supplied auto-reply policy in a new AUTOR00 parmlib member. This new function is expected to help provide a timely response to WTORs and help prevent delayed responses from causing system problems.
- XCF Status monitoring is designed to incorporate information about system-critical XCF group members that identify themselves, and initiate termination actions, including partitioning a system from the sysplex, if a monitored member fails to respond when polled for status or indicates impairment. This function is

- intended to help reduce the incidence of sysplex-wide problems that can result from unresponsive critical components. GRS exploits these XCF critical member functions in both ring and star modes. Additionally, GRS monitors key tasks and notifies XCF when it detects that GRS is impaired.
- z/OS V1.12, with z196 servers and Coupling Facility control code (CFCC) Level 17, is designed to capture Coupling Facility (CF) data nondisruptively in some circumstances, allowing the CF to continue operating. This new function is intended to help improve Parallel Sysplex availability when it is necessary to capture CF data.
 - SMF is updated to generate event-driven type 30 and type 89 records when a change in processor capacity is detected. Fields containing processor capacity data are added to the type 30 and type 89 records. This support also introduces a new SMFPRMxx parmlib option to govern the number of allowable event-driven intervals, and includes generating a new subtype of the type 90 record when a change in processor capacity is detected. The IFAURP billing report is also updated to allow for computing the SU factor from the data in the new type 89 records (values from a table are used when processing SMF records created before this function was available). Also, the reports are reformatted to structure the data by type and serial, rather than by type, model, and serial.
 - A new Predictive Failure Analysis check is designed to detect and automate the system's response to tasks that are writing SMF records at unusually high rates. Another new function, SMF record flooding automation, is designed to allow you to define a policy for responding to these situations in the SMFPRMxx member of parmlib, by specifying whether record flooding automation is to be active, whether operators are to be warned, and the actions to take for specific SMF record types if record flooding occurs. This is intended to limit the impact of such problems by allowing less-important data to be discarded while keeping the data from critical SMF records intact. Additionally, a new parameter is added to the SMF dump programs (IFASMF DL and IFASMF DP) to create a new flood statistics summary table you can use when developing a record flooding policy.
 - Predictive Failure Analysis (PFA) allows you to specify that it ignore data related to certain jobs or address spaces when you expect their behavior to be atypical of normal system operation. This action can help you improve the overall accuracy of PFA checks for LOGREC, message, and SMF record arrival rates as well as frames and slots usage. This action allows you to provide some inputs to the machine learning process, and is referred to as supervised learning. In z/OS V1.12, PFA will support both supervised and unsupervised learning. Also PFA will:
 - Capture data when exceptions are issued to help you identify problems
 - Use dynamic modeling intervals based on system stability
 - Provide improved common storage modeling
 - With z/OS V1.12 DFSORT:
 - A new DYNAPCT installation and run-time option is intended to improve reliability for the use of dynamically allocated work data sets.
 - DFSORT is designed to calculate more accurate file sizes for improperly closed VSAM fixed-length record data sets when possible.

- DFSORT is designed to issue diagnostic messages automatically in error situations, without the need to add a SORTDIAG DD statement or specify the DIAGSIM=YES installation option. This is intended to make it easier to provide the information you need to resolve many DFSORT problems without having to rerun a sort to collect diagnostic information, and to improve first failure data capture.
- z/OS UNIX System Services file system processing is designed to provide better information when a DISPLAY GRS,ANALYZE command is issued, by identifying itself as the holder of held latches to GRS. This is intended to help you diagnose and take corrective actions for latch contention problems that involve file system processing.
- Information about DFSMSrmm active and queued tasks is available via the DFSMSrmm API and via a TSO/E subcommand, enabling storage applications to monitor and act on the available information. In addition you can use the DFSMSrmm dialog to manage these tasks.
- In z/OS V1.12, RSM and dump processing are designed to improve capture performance for large amounts of data during SVC dump processing. This is expected to markedly reduce dump capture time when a large amount of data must be paged in during SVC dump processing. Internal IBM laboratory tests have shown that the capture time for SVC dumps can be reduced by up to 50-90% for large dumps with a substantial percentage of data on auxiliary storage. The amount of improvement is expected to vary depending on how much data must be paged in during dump processing, how much real storage is available to the system, and the system's workload.
- Dump processing is designed to act on a new option DEFERTND for the CHNGDUMP and DUMP commands, and in parmlib member IEADMCxx. The new option allows you to specify that task nondispatchability for address spaces being dumped as a result of a DUMP operator command be delayed until after global data capture is complete. This reduces the amount of time tasks and address spaces being dumped are set nondispatchable to capture volatile data to reduce the impact of command-initiated SVC dumps. This function is also available for z/OS V1.11 with the PTF for APAR OA29517.
- Stand-alone Dump is designed to better prioritize data capture for address spaces, and to dump a number of system address spaces first irrespective of their ASID numbers. This is intended to capture the data most often needed to diagnose system problems more quickly in case there is not enough time to take a complete stand-alone dump. Also, Stand-alone Dump is designed to allow you to specify additional address spaces to be added to the predefined list using a new ADDSUMM option.
- The existing XCF/XES CF structure hang detect support is extended by providing a new CFSTRHANGTIME SFM Policy option that allows you to specify how long CF structure connectors may have outstanding responses. When the time is exceeded, SFM is designed to drive corrective actions to try to resolve the hang condition. This helps you avoid sysplex-wide problems that can result from a CF structure that is waiting for timely responses from CF structure connectors.

- One focus area in z/OS V1.12 is the time it takes to restart the z/OS system itself and major subsystems such as DB2. Substantial reductions in restart times for DB2 systems that use a large number of data sets are expected in addition to improvements in the time required for some phases of z/OS initialization processing. Improvements include:
 - Design changes for Allocation, DFSMSdfp, and GRS intended to improve performance for address spaces that allocate a large number of data sets in a short time. These changes are expected to markedly reduce the startup time for these address spaces, such as DB2 and batch jobs that process a large number of data sets per job step.
 - Changing subsystem initialization from serial to parallel for initialization routines that are listed in IEFSSNxx parmlib members after a new BEGINPARALLEL keyword allows you to reduce system startup time by allowing many of these routines to run in parallel.
 - In z/OS V1.12, XCF has been changed to reduce the time it takes for the second and subsequent systems to join a Parallel Sysplex. This is intended to improve availability.
 - z/OS Communications Server now supports a hot-standby for sysplex distributor through the use of a new distribution method, HotStandby. You can configure a preferred server and one or more hot-standby servers. The preferred server with an active listener receives all new incoming connection requests, and the hot-standby servers act as backup should the designated preferred server become unavailable. The hot-standby servers can be ranked to control which server becomes the active server. You can also control whether the sysplex distributor automatically switches back to using the preferred server if it becomes available, and whether the distributor automatically switches servers if the active target is not healthy.
 - z/OS Communications Server sysplex problem detection and recovery is enhanced to detect when the TCP/IP stack has abended five times in less than a minute. Existing sysplex recovery logic is applied when this problem is detected.

Slide 29

Improving Availability 

z/OS UNIX System Services to use GRS latch identity

- New DISPLAY GRS,ANALYZE command introduced in z/OS R11
- Display easier to interpret for latch identity exploiters
- z/OS UNIX is a heavy user of latch services

SMF record flood detection, filtering

- New FLOOD option in SMFPRMxx and SETSMF
- Specify threshold and action to be taken in SMFPRMxx:
 - Issue console message
 - Suppress record and fail SMF(E)WTM request

Predictive Failure Analysis

- SMF record arrival rates
 - Too many of the same type records probably indicates a problem
 - Too few records probably indicates a different kind of problem
- Ability to specify atypical jobs and address spaces to be excluded from learning algorithms
- Data capture for exception conditions to help identify cause
- Modeling intervals based on system stability
- Discard LOGREC data from a period of time before a shutdown
- More-granular system AS common storage monitoring
- Common storage modeling now includes (E)SQA

Understanding GRS ENQ and Latch Usage and Contention Wednesday 3:00


z/OS PD Update—Wednesday 3:00

Detecting Soft Failures Using z/OS PFA: Tuesday 11:00



29  ©2010 IBM Corporation

Slide 30



Improving Availability 


Run Time Diagnostics

- New operator command to analyze key system indicators
- Goal is to help you identify the root of problems before system failure


Allocation, DFSMSdfp™, GRS performance

- Designed to improve performance for address spaces that allocate a large number of data sets
- Expected to markedly reduce the startup time for address spaces such as:
 - DB2 address spaces
 - Batch jobs that process a large number of data sets per job step



30  ©2010 IBM Corporation

Slide 31

Improving Availability 

Hot standby for Sysplex Distributor


- New HotStandby distribution method
- Can specify a preferred server and one or more hot standby servers
- Preferred server gets new incoming connection requests
- Hot standby servers serve when designated preferred server becomes unavailable
 - Can specify which hot standby server becomes active
- Can specify whether the system switches back to the preferred server if it becomes available and whether it switches servers if the active target is not healthy

What's New in z/OS Communications Server
Monday 3:00



31 ©2010 IBM Corporation

Slide 32

Improving Availability 

Structure-level recovery in Sysplex Failure Manager

- New CFSTRHANGTIME option to specify the time a structure can have outstanding requests
- SFM will drive actions against unresponsive connectors
 - Recovery
 - TerminationIntended to nip sympathy sickness problems in the bud


New XCF function to identify "critical members"

- New IXCJOIN option
- Member provides a function XCF can use to check on member health
- SFM to partition out systems with unresponsive critical members
- GRS exploitation

Nondisruptive CF Dumps

- With CFCC Level 17 on z196 servers
- Designed to capture CF data nondisruptively in some circumstances while the CF runs
- Intended to help improve Parallel Sysplex availability

z/OS Parallel Sysplex:
Update
Tuesday 11:00




32 ©2010 IBM Corporation

Slide 33


Improving Availability

SDUMP and Standalone Dump improvements

- New DEFERTND=(Yes/No) added on CHNGDUMP and DUMP commands, to specify task nondispatchability should be deferred until global capture is complete
- Standalone Dump to dump these system address spaces immediately after ASIDs 1-4, irrespective of their ASIDs:
 - ANTMAIN, CONSOLE, XCFAS, IOSAS, SMXC, WLM, CATALOG, GRS, SMF, ALLOCAS, ANTAS000, DEVMAN, DUMPSRV, GRS, IEFSCHAS, IXGLOGR, JESXCF, JES2, JES3, OMVS, PCAUTH, RASP, SMSPDSE, SMSPDSE1, SMSVSAM, TRACE
- SVC dump and RSM designed to reduce dump capture time
 - Significant improvement expected in auxiliary storage data capture
 - Dump exit exploiters can take advantage of improved RSM services to reduce the system impact of collecting large amounts of data
 - GRS dump exit now takes advantage of the improved RSM services
 - Internal tests show up to a 60% reduction in capture time
 - Based on internal IBM lab measurements

Smarter SVC Dump Processing for Improved z/OS Resiliency
Wednesday 11:00 

- Standalone Dump
 - Better prioritization of data capture for address spaces, some restartable system address spaces captured irrespective of their ASID numbers
 - You can specify additional address spaces to be dumped using a new ADDSUMM option

33  ©2010 IBM Corporation

Slide 34

Integrating new Applications and Supporting Industry and Open Standards

Language Environment support for calendar dates beyond 2038/2042:

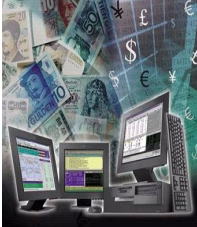
- `time_t` clock overflows in 2038 by UNIX architecture
- New `time64_t` service provides for creating calendar dates further in the future
- Available on z/OS R8 and up via APAR PK83212


Language Environment C/C++ support for BSAM Record I/O for >64K tracks

- DSNTYPE=LARGE support

z/OS R12 XL C/C++ Enhancements

- Metal C RENT support, including support for Metal C CICS transactions
- IPA(OBJECT) option behavior improvements
- More control over temporary data sets used during IPA link
- Implicit RESTRICT support, for deeper pointer analysis is intended to improve the performance
- Functions and parameters can be displayed for optimized code using optimization level O2 or O3 when you debug optimized code
- Many new compiler suboptions and macros, new pragma directives



34  ©2010 IBM Corporation

Application integration

The platform's classic strengths of availability, security, reliability, scalability, and management have made the mainframe the de facto standard for data serving and OLTP. It is logical to extend z/OS to Business Intelligence and Data Warehousing solutions as well, where large amounts of reports can be generated in a timely manner using source

data - all with a simplified reconciliation and restatement process. But it is also logical to deploy new or extend existing applications that leverage data on z/OS.

What sets z/OS apart is the ability to operate both new and existing applications within the same system, and in close proximity to your corporate data residing on z/OS. New applications based on Java, WebSphere Application Server, Perl, PHP, XML, C/C++, Unicode, HTML, HTTP, SOAP, z/OS UNIX System Services, and other Web services can operate side by side and integrate with classic applications based on CICS, IMS, DB2, Enterprise COBOL, Enterprise PL/I, REXX, System REXX, JCL, TSO/E, ISPF, Assembler, and Metal C. These applications can be colocated with relational (DB2) and non-relational (IMS) databases as well as record-oriented data. With such proximity to the data, applications on z/OS have a reduced need for expensive communications and networking infrastructure and can offer fewer opportunities for security breaches due to tight integration with traditional z/OS security, audit, and resource access.

Businesses with applications on z/OS understand the value of those applications, as well as understanding that replacing these systems with standard packages or other custom-built alternatives is not needed and may, in fact, be unnecessarily risky. Modernizing z/OS applications can lower costs and drive business agility with significantly enhanced levels of usability and integration.

z/OS V1.12 includes the following updates:

- Significant enhancements to z/OS XML System Services. eXtensible Markup Language (XML) is a W3C (World Wide Web Consortium) standard for the interchange of structured data, and is often considered a de facto standard for data interchange. For example, customers who have used DB2 for z/OS pureXML have found they have a much simpler and standard means of programmatically getting to DB2 for z/OS data, the ability to handle large volumes of large XML documents, and excellent price performance with z/OS XML System Services validating and non-validating parsing eligibility for the zAAP and zIIP specialty engines.
- For z/OS V1.12 z/OS XML System Services has performance improvements with new schema discovery and fragment parsing capabilities.
- C/C++ is one of the more popular application programming languages in the industry. For z/OS V1.12 and XL C/C++ V1.12 there are significant enhancements to performance, support of Euro currency, new standard time services, and Unicode.

Details on the application integration improvements for z/OS V1.12:

- z/OS XML System Services enhancements for z/OS V1.12 include:
 - z/OS XML System Services enhances XML schema validation support by allowing applications to extract schema location information from an XML instance document without the application first performing a separate parse operation. This improves the usability of the validating

parsing interface and reduces the processing cost of obtaining this information.

- z/OS XML System Services is updated to allow you to validate part of an XML document when performing validating parsing, rather than the entire document. Called fragment parsing, this capability reduces the processing cost of performing validation by allowing you to validate only a portion of a document rather than requiring the validation of the entire document. For example, this can be useful when only a subset of a large document containing multiple fragments has changed. The function is also planned to be made available for z/OS V1.10 and V1.11 with the PTF for APAR OA32251.
- z/OS XML System Services provides a new validating parse capability that allows applications to restrict the set of element names to be accepted as valid root elements to a subset of those allowable in an XML schema. This provides an additional level of validation capability beyond that provided by the W3C schema language. This function is also planned to be made available for z/OS V1.10 and V1.11 with the PTF for APAR OA32251.
- z/OS XML System Services provides new support designed to improve validating parsing performance for certain types of XML documents, in particular, those with complex schema, such as schema defined by the Single European Payments Area (SEPA) standard.
- SDSF makes Java classes available to provide a new means of accessing SDSF functions. Classes are provided for each SDSF panel for use by Java applications. This new support allows SDSF functions to be used easily from within Java-based applications.
- SDSF introduces a new ISFLOG command for SDSF REXX. It is designed to read the system log and return its records in stem variables with options to limit the number of records returned, and to allow you to specify start and end times. This new function simplifies access to the system log for SDSF REXX.
- Language Environment provides Euro currency support for Slovakia in the C/C++ Run-time Library. Both Euro and pre-Euro support are provided and the default locale for Slovakia will be changed to use the Euro symbol.
- Calendar times, represented by `time_t`, will overflow in January 2038. In z/OS V1.12, the Language Environment C/C++ Run-time Library includes new services including `time64_t`, which supports constructed calendar times up to and including 23:59:59 UTC on December 31, 9999.
- z/OS V1.12 Language Environment adds C99 compatibility support functions for C++ TR1 to bring the Standard C++ library in closer agreement with the library described in the ISO/IEC 9899:1999 Standard C, corrected through 2003:
- These C/C++ headers can now export C99-specific symbols, macros, functions, and function overloads: `complex.h`, `math.h`, `cctype`, `cmathcstdio`, `cstdlib`, `cwchar`
- C++ headers have been added: `cfenv`, `cinttypes`, `cstdbool`, `cstdint`
- Export symbols are now placed in the namespace `std::tr1`
- XL C/C++ delivers the following performance and usability enhancements for the z/OS V1.12 release:


- Metal C RENT support - The RENT option is enabled under the METAL option to support constructed reentrancy for Metal C programs with writable static and external variables, Metal C programs can now be concurrently used by multiple users, and Metal C can be used to write programs to run in CICS Transaction Server for z/OS, particularly in situations where you would otherwise code directly using Assembler.
- IPA(OBJECT) option behavior improvement - Starting with z/OS V1.12 XL C/C++, IPA(OBJECT) matches the behavior of IPA on other platforms, for consistency when using XL compilers on the other platforms. With IPA(OBJECT), object code is optimized by IPA-specific compile-time optimizations and in the IPA link step, IPA link-time optimizations are performed.
- New control over the temporary data set used during IPA link - A new IPA environment variable is now available to specify the size of the unnamed temporary work data set used during IPA link, allowing you to increase only the size of data set needed by IPA link instead of increasing all temporary work data sets.
- Implicit RESTRICT support - The RESTRICT option indicates to the compiler that parameter pointers in all functions or in specified functions are disjoint. When RESTRICT is in effect, deeper pointer analysis is done by the compiler, which is intended to improve the performance of the application being compiled.
- Message severity modification support - z/OS V1.12 XL C/C++ adds support for the SEVERITY option, to allow changing the severity of some of the diagnostic messages emitted by the compiler. This feature allows customization of certain diagnostic message levels to the coding standards of the user.
- Functions and parameters displayed for optimized code - For compiled code using optimization level O2 or O3, the names and addresses of functions and the names, types, and values of the parameters can be displayed when you debug optimized code.
- z/OS V1.12 XL C/C++ introduces the following new compiler suboptions:
 - INFO(ALS)
 - KEYWORD(typeof)for C source code
 - LANGLVL(AUTOTYPEDEDUCTION | C99LONGLONG | C99PREPROCESSOR | DECLTYPE | DELEGATINGCTORS | EXTENDEDINTEGERSAFE | INLINENAMESPACE | STATIC_ASSERT | VARARGMACROS | VARIADICTEMPLATES)
 - NAMEMANGLING(zOSV1R12_ANSI)
 - TARGET(zOSV1R12)
 - z/OS V1.12 XL C/C++ introduces the following new macros:
 - __IBMCPP_AUTO_TYPEDEDUCTION (C++ only)
 - __IBMCPP_C99_LONG_LONG (C++ only)
 - __IBMCPP_C99_PREPROCESSOR (C++ only)
 - __IBMCPP_DECLTYPE (C++ only)

- `__IBMCPP_DELEGATING_CTORS` (C++ only)
 - `__IBMCPP_INLINE_NAMESPACE` (C++ only)
 - `__IBMCPP_STATIC_ASSERT` (C++ only)
 - `__IBMCPP_VARIADIC_TEMPLATES` (C++ only)
 - `__PLO_INIT_PARAM_LIST`
 - `__PLO_PUTVAL_PARAM_LIST4`
 - `__PLO_PUTVAL_PARAM_LIST8`
 - `__PLO_PUTVAL_PARAM_LIST16`
 - `__PLO_GETVAL_PARAM_LIST8`
 - `__PLO_GETVAL_PARAM_LIST16`
- New pragma directives - z/OS V1R12 XL C/C++ introduces the new suboption `zosv1r12_ansi` to the option `NAMEMANGLING` and to the language construct `#pragma namemangling`. The linkage errors caused by any future changes in the mangling scheme can be resolved by recompiling the new modules using the same mangling scheme that was used in the already-compiled modules, `zosv1r12_ansi`.
- In z/OS V1.12, Program Management Binder supports new XPLINK versions to complement existing Binder C/C++ API DLL functions. These new XPLINK functions (`iewbnndx.so` and `iewbnndx.x`) are designed to help improve the performance of XPLINK applications by eliminating the overhead of XPLINK to non-XPLINK transitions when these bind functions are called. Also, a C/C++ header is provided to map the `IEWBMMP` structure (`__iew_modmap.h`). For C and C++ users, this will simplify the task of processing the module map the binder creates when the `MODMAP` option is used.
- There are a number of smaller Binder enhancements:
 - Sample programs that illustrate the use of both the Binder APIs and fast data access in High-Level Assembler and C
 - Character translations in `AMBLIST LISTLOAD` output for load modules
 - Improved `AMBLIST` header information for z/OS UNIX files
 - Support for long names for `AMBLIST LISTOBJ` for object modules in z/OS UNIX files
 - The Program Management Binder allows you to specify that a specific residency mode (`RMODE`) be applied to all initial load classes of a program object, rather than the classes in the first segment containing the entry point. This new function offers application programmers more flexible options for program storage residency.
 - The Program Management Binder makes program object attribute data (PMAR data) available to programs using fast data access, and fast data access now supports programs loaded using the z/OS UNIX System Services load service (`loadhfs`).
- Previously, the `tsocmd` shell command was available only from the Tools and Toys section of the z/OS UNIX System Services Web site. This function is now supported in z/OS V1.12. Unlike the existing `tso` command, the `tsocmd` command can be used to issue authorized TSO/E commands.
- Support is designed in z/OS UNIX System Services for the record file format in the `cp`, `mv`, `ls`, `pax`, and `extattr` shell commands as well as the `ISHELL` command.

- In addition to binary and text format, files can be handled in record file format. z/OS applications accessing these files by using QSAM, BSAM, VSAM, or BPAM and coding FILEDATA=RECORD will be able to take advantage of the record file format to read and write data as records.
- z/OS UNIX System Services now supports the memory mapping (mmap) function for files in zFS and HFS file systems. In z/OS V1.12, support allows applications to use memory mapping for NFS Client files. This will enable NFS-mounted file systems to be used by applications that use memory mapping.
 - A new option for the ISGENQ service can be used to serialize resources. This new support is designed to allow an unauthorized program to interrupt serialization processing and opt not to continue to attempt to obtain control of a resource when the resource is not available or to do other work asynchronously while waiting to obtain an ENQ resource. For example, a programmer might wish to set a time limit for obtaining control of a resource. This is expected to help programmers to better manage contention delays and remove pending enqueue requests in recovery.
 - JES2 and JES3 provide function you can use to specify, using the SYSOUT application programming interface (SAPI), that a program receive ENF 58 notifications when SYSOUT data sets have been deleted. This new function helps applications to monitor the progress of print data sets through the system.
 - The System Data Mover (SDM) component provides a REXX interface for many of the functions of the SDM programming interface (ANTRQST). This new function provides interfaces to FlashCopy®, Global Mirror (XRC), and Metro Mirror (PPRC) SDM services.
 - The CIM Server is upgraded to version 2.10 of the OpenPegasus CIM Server. Also, the CIM Servers Schema repository is updated to CIM Schema version 2.22. This is intended to keep the z/OS CIM Server and schema current with the CIM standard from OpenGroup and DMTF, and to allow z/OS enterprise management applications to manage z/OS systems. New CIM providers will support parts of the Host Discovered Resources (HDR) and Storage Host Bus Adapter (HBA) Profile from the SMI-S standard.
 - In z/OS V1.9 the C/C++ Run-time Library iconv() family of functions began to use Unicode Services to perform most character conversions. In z/OS V1.12, the ucmmap source or genxlt source for character conversions is removed from the C/C++ Run-time Library. You can create customized conversion tables using Unicode Services to replace these functions.
 - The WLM service for requesting LPAR-related data (REQLPDAT) is enhanced to include character-based data about the machine model, a Model-Permanent-Capacity Identifier, a Model-Temporary-Capacity Identifier, the Model-Capacity Rating, the Model-Permanent-Capacity Rating, and the Model-Temporary-Capacity Rating. This new data is intended to be used for reporting.
 - There are new functions for recovery and termination processing (RTM). These include a new option on SETFRR to specify that the SDWA may be above 16M, a new option on ESTAEX to specify that SPIE or ESPIE exits be superseded by EXTAEX, a new option on ESPIE to request percolation to RTM, and passing information about held locks to ESTAE-type recovery routines.

- A number of OPEN/CLOSE/End of volume (EOV) enhancements are made in z/OS V1.12:
 - A new DCB abend exit ignore option to additionally bypass the associated determinant abend message.
 - EOV processing will detect a missing last volume condition when reading, and if the JFCB and TIOT built by allocation have room, EOV will call catalog services to determine the additional volumes and then update the JFCB with the volume serial numbers and the TIOT with the device address. If this recovery is not possible EOV will abnormally terminate the job with abend 637-BC.
 - Externalize in a new message the reasons codes documenting the specific reason for a FREE=CLOSE failure.
 - Issue abend code 50D using a new reason code X'24'. Message IEC999I will continue to be issued with explanatory text.
 - Add new reason codes for abend 413 and abend 637 when detected for IEC709I, IEC710I, IEC711I, IEC712I. Also a recovery option will be made available to the DCB abend exit.

Slide 35


Integrating new Applications and Supporting Industry and Open Standards 

CIM Updates

- Storage provisioning CIM providers for SMI-S
 - Storage Management Initiative Specification owned by the Storage Networking Industry Association, also ANSI standard 388-2004 and ISO/IEC 24775
 - DMTF-based management based on CIM for storage
 - Support for mandatory providers for the Host Discovered Resources (HDR) and Host Bus Adapter (HBA) profiles
- CIM Server upgraded to a newer version of the OpenPegasus CIM Server
- CIM Servers Schema repository updated to CIM Schema version 2.22

WLM Request LPAR data (REQLPDAT) service to include:


- Character data for Machine model (e.g., x'F2F0F9F760F7F1F8' for a 2097-718)
- Character data for Model-Permanent-Capacity Identifier
- Character data for Model-Temporary-Capacity Identifier
- 32-bit unsigned Model-Capacity Rating
- 32-bit unsigned Model-Permanent-Capacity Rating
- 32-bit unsigned Model-Temporary-Capacity Rating



Intended to be used for reporting applications

35 ©2010 IBM Corporation

Slide 36

Integrating new Applications and Supporting Industry and Open Standards 

SDSF REXX™ Access to SYSLOG

- New ISFLOG command for SDFS REXX
- Designed to read SYSLOG and return records stem variables
- Options to limit the number of records returned, specify start and end times

SDSF support for Java and XML

- Allow easier access to SDFS functions from Java
- New Java classes, one per supported SDFS panel
- XML document input to SDFS, with XML document output
- SDSF XML schema shipped with z/OS
- Data and actions accessible from SDFS tabular panels available using XML

**SDSF Update
Thursday 11:00**


z/OS UNIX support for FILEDATA=RECORD

- Record types now supported include:
 - NA – Not defined
 - Binary – Binary
 - NL – Text delimited using newline characters
 - CR – Text delimited using carriage return characters
 - LF – Text delimited using line feed characters
 - CRLF – Text delimited using both carriage return and line feed characters
 - LFCR – Text delimited using both line feed and carriage return characters
 - CRNL – Text delimited using both carriage return and newline characters
 - Record – Prefixed records in which the prefix defines the record length**

**z/OS V1.12 UNIX System
Services Latest Status and New
Features
Monday 4:30**

36 ©2010 IBM Corporation

Slide 37


Integrating new Applications and Supporting Industry and Open Standards 

z/OS UNIX System Services support for mmap for NFS files


- Designed to provide memory mapping for NFS Client files

z/OS XML System Services enhancements

- Allow applications to extract schema location information from an XML instance document without the application first performing a separate parse
 - More usable, expected to have lower processing cost
- Allow validation of part of an XML document when performing validating parsing
 - Fragment parsing is intended to offer lower processing cost during validation by allowing you to validate only a portion of a document
- New validating parse capability to allow applications to restrict the set of element names to be accepted as valid root elements to a subset of those allowable in an XML schema
 - Intended to provide another level of validation, beyond W3C schema


37  ©2010 IBM Corporation

Slide 38


Integrating new Applications and Supporting Industry and Open Standards 

Binder Improvements:

- XPLINK version of the Binder C/C++ API DLL functions
 - Designed to offer XPLINK applications improved performance by eliminating expensive XPLINK to non-XPLINK transitions
 - New C/C++ header to map the IEWBMMP structure, to make it simpler to process module map created by MODMAP
- Sample HLASM and C programs for standard and Fastdata Binder APIs
- Character translations in AMBLIST LISTLOAD output for load modules
- Improved AMBLIST header information for z/OS UNIX files
- Support for long names for AMBLIST LISTOBJ for object modules in z/OS UNIX files
- Allow you to specify an RMODE to be applied to all initial load classes of a program object
- Make program object attribute data (PMAR data) available to programs using the fast data interface, and to support programs loaded using the z/OS UNIX System Services load service (loadhfs)


38  ©2010 IBM Corporation

Slide 39

Integrating new Applications and Supporting Industry and Open Standards 


New ISGENQ Service Option

- Allow unauthorized programs to interrupt serialization processing and choose whether to:
 - Continue to obtain control of a resource
 - Not to continue to attempt to obtain control of a resource
 - Do other work asynchronously while waiting
- For example, you might wish to set a time limit for getting control of a resource
- Expected to help programmers to better manage contention




System Data Mover (SDM)

- New REXX interface for many ANTRQST functions
- Designed to provide interfaces to FlashCopy®, Global Mirror (XRC), and Metro Mirror (PPRC) SDM services


39  ©2010 IBM Corporation


Slide 40

Enhancing Security 

IBM Tivoli® Directory Server (LDAP) Improvements

- **Password policy**
 - Lockout
 - Password quality
 - Expiration
 - Password reuse ...
- **ACL extensions to provide access control based on:**
 - IP address
 - Distinguished name (DN)
 - Time of day and week
 - Certain bind characteristics
- **Activity log management**
 - Set maximum size
 - Archive or delete inactive log files
- **Salted SHA for passwords**
 - Help make dictionary attacks more difficult
- **Support for other IBM Tivoli Directory Servers**
 - Designed to allow migration and replication of schema and directory entries from IBM Tivoli Directory Server instances on other platforms



40  ©2010 IBM Corporation

Security

Security is often a moving target. New security-related capabilities are often followed by ever-more sophisticated and creative attempts to circumvent them. z/OS has a huge breadth of security capabilities built into the base of the operating system at no extra cost. Many z/OS security functions, such as data encryption, encryption key management,

digital certificates, password synchronization, and centralized authentication and auditing, can be deployed as part of enterprise-wide security solutions and can help mitigate risk and reduce compliance costs, while accelerating and reducing cost of implementation.

Encryption obscures information and is intended to make it unreadable to unauthorized parties. Cryptography can be used to protect the confidentiality, integrity, and availability of both data at rest and data being transmitted, and in general remains one of the strongest aspects of IT security. z/OS is the logical choice for cryptography and storing and managing the cryptographic keys due to the nature of key handling by z/OS Integrated Cryptographic Service Facility (ICSF). The Integrated Cryptographic Service Facility could be considered more secure than other cryptographic solutions because it can manage the encryption and decryption of sensitive material without exposing the keys in clear. z/OS V1.12 is updated with many cryptographic capabilities, such as: support for new smart card formats, new cryptography standards and algorithms (such as DSA, DH, EC, BLOWFISH, RC4, Galois/Counter Mode encryption for AES (GCM), Elliptic Curve Cryptography (ECC), Elliptic Curve Diffie-Hellman key derivation (ECDH), Elliptic Curve Digital Signature Algorithm (ECDSA), and Hashed Message Authentication Mode (HMAC)), as well as z/OS Communications Server support for IKEv2 and Federal Information Processing Standard (FIPS) FIPS 140-2.

z/OS meets many of the industry's toughest security standards. For z/OS V1.12, the z/OS Communications Server has new support for IKEv2 and Federal Information Processing Standard (FIPS) FIPS 140-2. In addition, in June 2010, IBM z/OS V1.11 was certified as meeting the requirements of the recently published German Common Criteria Certification Body (BSI) Operating System Protection Profile (OSPP). z/OS Version 1 Release 11 is the first operating system to be certified as meeting the requirements of this new Protection Profile.

Digital certificates are used in managing and working with private key/public key encryption and are often required as part of security and compliance guidelines. They can be used by applications to establish secure communication sessions or to configure virtual private network (VPN) sessions, and to authenticate users and objects. z/OS PKI Services is a complete digital certificate authority included in the base of z/OS at no additional charge. Relatively few z/OS resources can be used to generate hundreds of thousands, even millions of digital certificates. Reduce risk and reduce cost by generating and managing your own digital certificates from z/OS. For z/OS V1.12 z/OS PKI Services is enhanced with several usability improvements which are expected to reduce the amount of time and number of manual tasks associated with finding certificate serial numbers, and issuing renewal and revocation e-mails. Support for additional standards, such as Certificate Management Protocol (CMP), means devices can now request, revoke, suspend, and resume certificates from z/OS PKI Services directly and automatically. Certificates generated by z/OS PKI Services can also be customized for use with Microsoft Exchange and smart card readers.

Authentication, auditing, and compliance are growing concerns. Many laws and standards have been recently refined, enacted, or created, governing the protection and access of data. z/OS has a long history of resource access and reporting capabilities built into the platform that can be useful for administering z/OS security, monitoring for threats, and auditing usage and policy compliance. z/OS V1.12 has significant updates for IBM Tivoli Directory Server for z/OS (LDAP) in support of new password policy rules, improved activity logging, and new extensions for access control lists.

Details on the security enhancements for z/OS V1.12:

- ICSF provides support for translation of external RSA tokens wrapped with key encrypting keys into one of three smart card formats. A new callable service, PKA Key Translate (CSNDPKT), is designed to translate an existing RSA private key in CCA external format into a specified smart card (SC) format in support of VISA, or the common ME or CRT format. To use this new function, you will need a zEnterprise 196 server, or an IBM System z10 or System z9 server with the Crypto Express2 feature with minimum driver and microcode levels dated April 2009 or later. This function is also available on z/OS V1.10 with the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable and PTF UA46713 or with z/OS V1.11 or later.
- The Central Processor Assist to Cryptographic Function (CPACF) on zEnterprise 196 and IBM System z10 servers with the CEX3C feature is designed to help facilitate the continued privacy of cryptographic key material when used by the CPACF for high-performance data encryption. Leveraging the unique z/Architecture®, protected key CPACF is designed to help ensure that key material is not visible to applications or operating systems when used for encryption operations. Protected key CPACF is designed to provide significant throughput improvements for large volumes of data and low latency for small blocks of data. In z/OS V1.12, ICSF exploits the enhancements made to the CPACF in support of separate key wrapping keys for DES/TDES and AES. This is designed to provide the same functions available using the PCI card, but with the advantage of CPACF performance.
- There are a number of improvements for PKI Services in z/OS V1.12.
 - PKI Services allows you to create and sign certificates with Elliptic Curve Cryptography (ECC) keys, in addition to RSA and DSA keys. See more information on ECC below.
 - PKI Services support longer distinguished names in digital certificates. This is intended to support the generation and use of certificates with very long distinguished names.
 - Support is provided for the creation of certificates with expiration dates in the far future to give greater flexibility on certificate validity period for customers.
 - Automatic adjustment on serial number to be issued avoids problems caused by duplicated serial numbers.
 - PKI Services provides utilities to post certificates and Certificate Revocation Lists (CRLs) to LDAP on demand.

- Certain tasks, such as removing old or expired certificates and requests, and processing certificate expiration notification warning messages, have historically consumed considerable processing time when you have a large number of certificates. A new PKI Services design is intended to markedly improve the performance and reduce the processing time of these tasks and allow you to specify the time of day and days of the week this task will be run.
- PKI Services is designed to support custom extensions to X.509 Version 3 certificates, for example, an extension called Certificate Template Name for Microsoft Exchange or Smart Card Login.
- PKI Services allows you to create a certificate with Subject Alternate Name that contains multiple instances of each of the supported General Name forms.
- Certificate Management Protocol (CMP) is an Internet protocol used to manage X.509 digital certificates described by RFC 4210, which uses the Certificate Request Message Format (CRMF) described by RFC 4211. In z/OS V1.12, PKI Services provides support for parts of the CMP standard, allowing CMP clients to communicate with PKI Services to request, revoke, suspend, and resume certificates. This is intended to allow you to use CMP in a centralized certificate generation model.
- PKI Services supports passing the reason a certificate request was rejected from the administrator to the requester in the rejection e-mail.
- RACDCERT enhancements include:
 - Support for the creation of certificates with expiration dates in the far future to give greater flexibility on certificate validity period for customers.
 - Support for longer distinguished names in digital certificates. This is intended to support the generation and use of certificates with very long distinguished names.
 - Support for Elliptic Curve Cryptography (ECC). See more information below.
 - In 2009, the U.S. National Institute of Standards and Technology (NIST) published an IPv6 profile that requires support of certain cryptographic suites as defined in RFC 4869, Suite B Cryptography Suites for IPsec. One of the technologies referenced was Elliptic Curve Cryptography (ECC), which is regarded as providing stronger cryptography with smaller key sizes than RSA cryptography. This type of cryptography is expected to be attractive for use with small devices such as mobile devices and smart cards, that have limited computing power. In z/OS V1.12, PKI Services allow you to create and sign certificates with ECC keys in addition to RSA keys. In z/OS V1.12, System SSL provides support for ECC-related data structures, signing data, and verifying signed data using ECDSA (Elliptic Curve Digital Signature Algorithm). This is intended to allow exploiters of z/OS System SSL to import ECC style certificates and private keys into key database files or PKCS#11 tokens and use ECDSA certificates in signing and verifying operations. In z/OS V1.12, the RACF RACDCERT

command allows you to create and sign certificates with ECC keys, in addition to RSA and DSA keys.

- A discrete general resource profile with generic characters (*,%,&) in its name, defined in a class enabled for generics (GENCMD or GENERIC), is often called a "ghost" profile. Such profiles are not referenced by RACF for authorization checking. However, when defined, they can confuse and annoy RACF administrators and system programmers. In z/OS V1.12, RACF provides a new NOGENERIC keyword for the RDELETE command to enable you to delete these profiles along with a GENERIC=N option for R_admin DELETE.
- The Command Prefix Facility (CPF), which you can use to route commands from one system to another within a sysplex, supports security checking similar to that provided for the ROUTE operator command. Defining a new MVS™.CPF.ROUTE.CHECK profile in the RACF OPERCMDS class specifies that the system use the MVS.ROUTE.CMD profile in the RACF OPERCMDS class to determine whether the operator is allowed to send a command to the specified system. This is intended to add the same level of checking to CPF that exists for the MVS ROUTE command.
- The Network Authentication Service for z/OS utilizes RACF function to help improve the availability of applications that use Kerberos or GSSAPI services when deployed in a DVIPA environment. This new support is designed to allow you to remove the dependency on which image of the Sysplex a Kerberos or GSSAPI application request is routed to. This can help improve application availability by enabling transparent failover and improved workload balancing between images in a Sysplex. This function is also available for z/OS V1.11 with the PTF for APAR OA32352.
- IBM Tivoli Directory Server for z/OS (LDAP) provides a number of new enhancements:
 - IBM Tivoli Directory Server for z/OS provides support for configurable password policy rules that can be applied to user passwords in the directory. Support for automatic password revocation, password expiration, formatting checks, history, and a password change mechanism can be enforced on an individual, group, or directory basis. This new function is intended to help you ensure that:
 - Users change their passwords periodically
 - New passwords meet your password requirements
 - Recently used passwords not be reused
 - Users can be locked out after a defined number of failed attempts
 - In addition, when a password policy control has been received, native or SDBM authentication will map RACF response codes to password policy response codes where possible, and the password policy response codes will be returned.
 - IBM Tivoli Directory Server for z/OS supports continuous activity logging. This new function is designed to close the current log file or generation data set and open a new one based on the time of day or the size of an activity log file you specify. The console command is designed to allow initiation of an activity log file switch. Also in this support is a

new function that will allow you to specify that log entries are to be filtered by IP address.

- IBM Tivoli Directory Server for z/OS provides an extension to access control lists (ACLs) to provide the ability to dynamically transform base ACLs using ACL filters you specify to add or remove permissions based on:
 - Bind distinguished name (DN)
 - Alternate DNs
 - Pseudo DNs
 - Groups a bind or alternate DN belongs to
 - IP address of the client connection
 - Time of day that directory entry was accessed
 - Day of week that directory entry was accessed
 - The bind mechanism used
 - Whether bind encryption was used

This function is designed to provide additional flexibility in access controls for LDAP connections.

- IBM Tivoli Directory Server for z/OS provides Salted SHA-1 encryption support. Intended to make dictionary attacks against SHA-1 encrypted data much more difficult, stored Salted SHA-1 password values in LDAP will include a random 20-byte string so that encrypting the same password more than once will usually result in differing encrypted values. This is intended to make it much more difficult to determine the encrypted password value. This support is designed to be functionally equivalent to that currently provided by the IBM Tivoli Directory Server and can allow easier migration of LDAP server workloads to z/OS.
- IBM Tivoli Directory Server for z/OS provides support for the syntaxes and matching rules currently supported by IBM Tivoli Directory Server. This support is designed to allow migration and replication of schema and directory entries using these syntaxes and matching rules from IBM Tivoli Directory Server on other platforms.
- The NFS Server supports password phrases up to 100 characters in length for mvslogin, in addition to existing support for passwords up to 8 characters long. This function requires password phrase support from RACF, or another external security manager. This is intended to allow you to migrate to password phrases, which offer a much larger name space than passwords.
- TSO/E now supports passwords that include one or more special characters. This is intended to leave the checking for acceptable password characters to an external security manager such as RACF.
- z/OS Communications Server introduces trusted TCP connections, to enable sockets programs to retrieve sysplex-specific connection routing information and partner security credentials for connected sockets. Partner security credentials can be retrieved if both endpoints of a TCP connection reside in the same z/OS image, z/OS sysplex, or z/OS subplex, and the endpoints are within the same security domain. In such a topology, partner programs can use trusted connections to

authenticate each other as an alternative to using an SSL/TLS connection with digital certificates for client and server authentication.

- Internet Key Exchange version 2 (IKEv2) is the latest version of the Internet Key Exchange (IKE) protocol specified by RFC 4306. IKE is used by peer nodes to perform mutual authentication and to establish and maintain security associations (SAs). In z/OS V1.12 the Communications Server IKE daemon (IKED) is enhanced to support IKEv2, in addition to its existing IKEv1 support. The z/OS Communications Server support for IKEv2 includes:
 - IPv4 and IPv6 support
 - A new identity type called KeyId
 - Authentication using pre-shared keys or digital certificates; certificates may use RSA or elliptic curve (ECDSA) keys
 - Re-keying and re-authentication of IKE SAs and child SAs
 - Hash and URL specification of certificates and certificate bundles
 - A new certbundle command which can create certificate bundles as specified by RFC 4306
- z/OS Communications Server introduces these enhancements to the network security services daemon (NSSD) IPsec Certificate Services:
 - IKEv2 support: X.509 certificate-based signature creation and validation for IKEv2.
 - Elliptic Curve Digital Signature Algorithm (ECDSA) support: X.509 certificates that contain ECDSA keys may be utilized for IKEv2 digital signature creation and verification.
 - X.509 certificate trust chain support: The entire X.509 trust chain will be taken into consideration during IKEv1 or IKEv2 digital signature creation and verification.
 - Certificate Revocation List (CRL) support: CRLs may be retrieved via HTTP and consulted during IKEv1 or IKEv2 digital signature verification.
 - Hash and URL support: Certificates and certificate bundles specified using the Hash and URL format specified in RFC 4306 may be utilized during IKEv2 digital signature creation and verification.

The z/OS Internet Key Exchange daemon (IKED) is enhanced to use these new NSSD functions when a stack is configured as a network security client.
- z/OS Communications Server introduces these enhancements to IPsec and IKE support for cryptographic currency:
 - Support for the Advanced Encryption Standard (AES) algorithm in Cipher Block Chaining (CBC) using 256-bit keys, an addition to the previously existing 128-bit key support. You can use the longer key length for more sensitive data.
 - Support for the Advanced Encryption Standard (AES) algorithm in Galois Counter Mode (GCM) and in Galois Message Authentication Code (GMAC) mode. AES in GCM is intended to provide both confidentiality and data origin authentication. AES-GCM is regarded as a very efficient algorithm for high-speed packet networks. AES in GMAC mode is intended to provide data origin authentication but does not provide

confidentiality. AES-GMAC, like AES-GCM, is also regarded as a very efficient algorithm for high-speed packet networks. z/OS V1.12 Communications Server supports both 128-bit and 256-bit key lengths for these algorithms.


- Support for the use of Hashed Message Authentication Mode (HMAC) in conjunction with the SHA-256, SHA-384, and SHA-512 algorithms. These algorithms are intended to be used as the basis for data origin authentication and integrity verification. The new algorithms, HMAC-SHA-256-128, HMAC-SHA-384-192, and HMAC-SHA-512-256, are designed to help ensure that data is authentic and has not been modified in transit. Versions of these algorithms that are not truncated are available as Pseudo-Random Functions (PRFs). These algorithms are called PRF-HMAC-SHA-256, PRF-HMAC-SHA-384, and PRF-HMAC-SHA-512.
- Support for an authentication algorithm, AES128-XCBC-96, that can help ensure data is authentic and not modified in transit.
- Support for Elliptic Curve Digital Signature Algorithm (ECDSA) authentication.
 - Support for Elliptic Curve Diffie-Hellman (ECDH) key agreement.
- z/OS Communications Server IPsec and IKE support leverages z/OS cryptographic modules that are designed to address the Federal Information Processing Standard (FIPS) FIPS 140-2 security requirements for cryptographic modules. FIPS 140 defines a set of security requirements for cryptographic modules to obtain higher degrees of assurance regarding the integrity of those modules. FIPS 140-2 provides four increasing, qualitative levels of security intended to cover a wide range of potential applications and environments. z/OS V1.12 Communications Server support is designed to be configurable such that it will only utilize underlying security modules (System SSL and ICSF's PKCS #11 capabilities) designed to address the requirements for FIPS 140-2 level 1.
- RFC 4301 compliance for z/OS Communications Server IPsec filter rules becomes mandatory. RFC 4301 "Security Architecture for the Internet Protocol" specifies the base architecture for IPsec-compliant systems, including restrictions on the routing of fragmented packets. Compliance enforcement may require minor changes to IP filters for IP traffic that is routed through z/OS. IBM Configuration Assistant for z/OS Communications Server is designed to identify any non-compliant IP filters and the policy agent is designed not to install an IPsec policy that contains any non-compliant IP filters.
- In prior releases, System SSL supported X.509 certificates with RSA key sizes up to 2048 bits for use in PKCS #11 tokens. In V1.12, System SSL gskkyman is enhanced to support the creation and management of X.509 certificates and keys within a PKCS #11 token that have RSA key sizes up to 4096 bits, DSA keys and Diffie-Hellman keys. These X.509 certificates and keys are intended to be usable through the System SSL APIs.
- The Cryptographic Support for z/OS V1.10 through z/OS V1.12 Web deliverable 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, and include these new functions:
 - 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 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 Express® 3 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 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.

Slide 41


Enhancing Security 

ICSF*


ICSF Update
Thursday 9:30

- New PKA Key Translate service designed to translate CCA RSA private keys into a VISA smart card (SC) format or the common ME or CRT format
 - Requires a z196, IBM System z10™, or IBM System z9® server with the Crypto Express2 feature, minimum driver and microcode levels dated April 2009 or later, z/OS V1.11, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable and PTF UA46713
- Cryptographic Support for z/OS V1.10 through z/OS V1.12 Web deliverable is planned to be made available September 10, 2010*, is designed to support:
 - Significantly more keys in a CKDS and reduced update overhead
 - 64-Bit enablement for the full set of ICSF callable services
 - Identity information in ICSF SMF Type 82 audit records
 - API access to new symmetric encryption modes introduced in 196 server's CPACF, including CMAC, CCM, GCM, CFB, OFB, XTS and CBC-S for both CPACF clear keys and CPACF protected keys
 - New functions on z196 servers with the Crypto Express3 feature, with minimum driver and microcode levels dated August 2010 or later, and the PTF for APAR OA33260:
 - CCA key token wrapping designed to meet ANSI x9.24-1 key management requirements for symmetric keys wrapped under symmetric master or key-encrypting
 - Enhanced PIN protection for verification at ATM and POS systems intended to provide defined extensions to ANSI x9.8 and ISO 9564-1 standards
 - New CCA functions implementing ECDSA, including support for NIST and Brainpool standard curves, and for the Crypto Express 3 Coprocessor performing ECDSA within the bounds of the secure hardware
 - New CCA HMAC functions for SHA-1 and SHA-2

41 * 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



Slide 42

Enhancing Security* 

Elliptical Curve Cryptography support in RACF, PKI Services, and System SSL

- SHA224, SHA256, SHA384, SHA512 support
- RACDCERT and PKI certificate support based on draft RFC3279

SAF controls on sysplex-scope CPF commands

- New authorization check for appropriate MVS.ROUTE.CMD.* profile in the OPERCMDS class when you define a new profile, MVS.CPF.ROUTE.CHECK

Digital Certificate support for:


- Prevention of duplicate certificate serial numbers
- Certificate Management Protocol (CMP) in PKI Services to allow clients to request certificate operations

RACDCERT and PKI support for really long certificate expiration dates and distinguished names

- Dates previously limited to 2038 (PKI) and 2041 (RACDCERT)
- Expiration dates supported up to "only" the year 9999
- DN length limit changed from 255 to 1024
- Available on z/OS R10 and up with PTFs for APARs OA30560 and OA30952


Understanding Digital Certificates on z/OS **RACF Update**
Wednesday 11:00 **Wednesday 4:30**

42 ©2010 IBM Corporation




Slide 43

Enhancing Security



Additional PKI Services enhancements:


- New utility to allow you to post existing certificates in LDAP automatically
- New utility to allow you to post updates to CRLs immediately
- Designs to improve performance for "housekeeping tasks"
 - Removing old and expired certificates and requests
 - Processing certificate expiration notification warning messages
- Ability to specify the time of day and days of the week for housekeeping tasks
- Support for passing the reason a certificate request was rejected from the administrator to the requester in the rejection e-mail
- Custom extensions to X.509 Version 3 certificates, such as:
 - Creating Domain Controller certificates with Certificate Template Name extensions, with an OIDs, and with BMP data "DomainController" for use with Microsoft Exchange or Smart Card Login
- Support for creating certificates with Subject Alternate Names containing multiple instances of each of the General Name forms support
 - For example, specifying multiple IP addresses when one was previously allowed



43 ©2010 IBM Corporation

Slide 44

Self-Managing Capabilities*




Capacity Provisioning

- Capacity Provisioning support for an "averaged rolling" Performance Interval
 - Prior support based on *n* consecutive intervals in which goals were not met
 - Now, can specify that if goals have not been met for most of the intervals CP should consider a provisioning action
 - Available for z/OS R10 and up with the PTF for APAR OA30496
- CP monitoring of CICS and IMS transaction classes
 - Currently, service classes for transactions have zero CPU delay counts
 - New "Total delay count" indicates whether transactions managed this way are delayed due to CPU, so provisioning can be driven if needed
 - Available for z/OS R10 and up with the PTF for APAR OA29641

**Introduction to z/OS Capacity Provisioning
Wednesday 4:30**

- Command Correlation
 - Help tie command completion events to CP requests
 - Available for z/OS R10 and R11 with the PTF for APAR OA30496



44 ©2010 IBM Corporation

Optimization and management capabilities

With the ability to intelligently manage workloads, reprioritize work, dynamically reallocate system resources between applications quickly and efficiently, and help meet business priorities, z/OS and System z can handle unexpected workload spikes and help improve your system's efficiency and availability.

New real-time RMF-based workload reporting. The z/OS Management Facility V1.12 (5655-S28) Sysplex Status and Monitoring Desktops applications can provide real-time status of resources of all your servers, sysplexes, and Linux images from one location. New designs provide for cross-sysplex performance monitoring from a single point with a quick red-yellow-green health indicator for your systems on a single panel, and it takes just seconds to see the health of all your sysplexes (and Linux for System z and Linux for Intel images). For example, if a system within a sysplex is missing goals, you can get a fast drill down on system components using the Monitoring Desktops application in as few as three clicks to view real-time statistics.

DFSMSdss and DFSMSHsm exploit the Fast Reverse Restore feature of the IBM System Storage DS8000™ Series. This function allows recovery to be performed from an active, original FlashCopy target volume to its original source volume without having to wait for the background copy to finish when the volume pair is in a full-volume FlashCopy relationship.

The z/OS Workload Manager (WLM) is a cornerstone to z/OS leadership in on demand computing. WLM lets you define performance goals and assign priorities to the goals. You can define the goals for work in business terms, and let the system decide what resources, such as CPU and storage, should be given to the workloads it to meet each goal. Workload Manager monitors the system and adapts processing to try to meet the goals. The scope of the Workload Manager extends to many aspects of system processing, from helping the management of incoming TCP/IP and SNA traffic to managing requests for I/O. You can take advantage of WLM to manage workloads using z/OS middleware, like DB2, CICS, IMS, WebSphere MQ, and other WebSphere products, and to manage the priority and execution of transaction requests across your z/OS Systems.

In addition to WLM enhancements, batch processing windows can be shortened and optimized through several other z/OS capabilities. z/OS V1.12 Catalog management has been updated to allow system functions including the IDCAMS utility to avoid DFSMSHsm recalls when deleting generation data groups (GDGs). As a reminder, in z/OS V1.11, job streams using the IEFBR14 program can be run faster by enabling Allocation to delete migrated data sets without first recalling them.

Extending the scope of z/OS and System z management, the Capacity Provisioning Manager for z/OS enables z/OS and the System z10 and later servers to add temporary capacity automatically when necessary, with or without operator intervention. Capacity Provisioning for z/OS V1.12 has been enhanced to use CICS and IMS monitoring data to determine whether additional resources are needed to meet service level requirements for these workloads. What has taken minutes or hours to discover, identify, decide, and resolve, now can be specified to happen automatically in one to two minutes. This function is also available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA29641.

Details on the optimization improvements for z/OS V1.12 include:

- The creation of new VSAM data sets with IMBED and REPLICATE attributes has been unsupported since z/OS V1.3. These attributes, originally introduced to improve performance on older DASD, typically act only to occupy additional space and slow performance on modern cached DASD. In z/OS V1.12, the system is designed to remove these attributes automatically from VSAM data sets logically dumped using DFSMSdss and migrated using DFSMSHsm when DFSMSdss is used as the data mover during restore and recall processing. An informational message will be issued to confirm that newly restored data sets no longer retain these attributes.
- DFSMSdss and DFSMSHsm exploit the Fast Reverse Restore feature of the IBM System Storage DS8000 Series. This function allows recovery to be performed from an active, original FlashCopy target volume to its original source volume without having to wait for the background copy to finish when the volume pair is in a full-volume FlashCopy relationship. DFSMSdss has been enhanced to create full-volume copies using a new keyword in order to support a Fast Reverse Restore function. DFSMSHsm FlashCopy backup and recovery operations have been enhanced to create full-volume FlashCopy relationships when the devices support it. The Fast Reverse Restore function supports the recovery of volumes associated with copy pool backups including Space Efficient and Incremental FlashCopy targets. A new DFSMSHsm SETSYS parameter allows you to specify whether extent or full-volume FlashCopy relationships are to be established between volume pairs when DFSMSHsm invokes DFSMSdss to perform fast replication backup and recovery.
- DFSMSrmm helps with reporting of data sets and logical volumes which are copy exported from a TS7700 virtualization engine.
- z/OS V1.12 has improved batch management with the following:
 - WLM considers resource group maximums and the projected increase in system or sysplex demand before starting initiators during resource adjustment and policy adjustment processing when the service classes have been assigned to resource groups and resource group maximums have been defined. The Type 99 SMF record is extended to show when the number of initiators to be started was limited for this reason. These changes are intended to improve WLM batch management.
 - A new parmlib option in IEAOPTxx allows you to specify that the system run CPU-intensive discretionary work for a longer period of time before dispatching other discretionary work, while still interrupting it after short periods for nondiscretionary work. This change is intended to help improve the throughput for systems with CPU-intensive discretionary workloads.
 - Catalog Management is enhanced to avoid DFSMSHsm recalls for any generation data sets that are migrated when deleting entire generation data groups (GDGs). Instead, Catalog Management will call DFSMSHsm to delete such data sets without recalling them. This is expected to reduce

processing time, particularly when one or more generation data sets have been migrated to tape.

- Initiator address spaces consume processor time on behalf of starting and ending job steps that in prior releases is not associated with a particular batch job. There can be considerable variation in the processor time consumed by an initiator for different jobs. To help you better understand the resources consumed by batch jobs and improve the accuracy of chargeback programs, z/OS V1.12 is designed to record the CPU time consumed in initiator address spaces using new fields in SMF Type 30 records.
- Capacity Provisioning uses the delay data for transaction service classes provided by RMF to help determine whether a provisioning action is required for servers on which CICS and IMS are running. Monitoring delay data for CICS and IMS transaction classes is intended to help improve capacity provisioning decisions for servers with LPARs running CICS and IMS. This function is also available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA29641.
- The Capacity Provisioning Manager allows you to specify that it is to use performance index rolling averages to determine whether a provisioning action should be taken rather than fixed intervals. This is intended to help improve the responsiveness of capacity provisioning. This function is also available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA30496.
- In z/OS V1.12, z/OS Communications Server uses new TCP/IP callable NMI requests to provide TCP/IP stack network interface information and network interface and global statistics. Network management applications can use the requested output to monitor interface status and TCP/IP stack activity. z/OS V1.12 Communications Server support is provided for these new requests:
 - GetGlobalStats - Provides TCP/IP stack global counters for IP, ICMP, TCP, and UDP processing
 - GetIfs - Provides TCP/IP network interface attributes and IP addresses
 - GetIfStats - Provides TCP/IP network interface counters
 - GetIfStatsExtended - Provides data link control (DLC) network interface counters
 - z/OS Communications Server provides enhancements to help you improve the management of the CSSMTP application and provides additional DVIPA-related sysplex event information in new subtypes of SMF Type 119 records:
 - DVIPA status change (subtype 32)
 - DVIPA removed (subtype 33)
 - DVIPA target added (subtype 34)
 - DVIPA target removed (subtype 35)
 - DVIPA target server started (subtype 36)
 - DVIPA target server ended (subtype 37)
 - CSSMTP Configuration data records (subtype 048)
 - CSSMTP Target server connection records (subtype 049)
 - CSSMTP Mail records (subtype 050)

- CSSMTP Spool records (subtype 051)
- CSSMTP Statistics records (subtype 052)


Applications can obtain the new SMF 119 subtypes from an SMF exit routine or in real time from the z/OS Communications Server Network Management Interface (NMI) for SMF, SYSTCPSM.

- CSSMTP issues the SIOCSAPPLDATA ioctl to add application data (appldata) to the TCP connections used to connect to target mail servers. You can see the application data (appldata) displayed in the Netstat All/-A, AllConn/-a, and Conn/-c reports.
- z/OS V1.12 includes several Network File System (NFS) enhancements:
 - NFS Server now provides support for SMF records to be created when file system objects are created, renamed, or removed. These new SMF Type 42 subtype 26 records can be specified with a new smf site attribute or enabled by the MODIFY operator command for z/OS UNIX System Services file objects including directories and files, and also for data sets and members when accessed using NFS. In conjunction with existing z/OS UNIX Type 92 and RACF Type 80 SMF records, this support is intended to make it easier to audit NFS file and data set operations.
 - The NFS MVSLOGIN client utility and NFS Server now support the use of password phrases from 14 to 100 characters long when SAF is used for RPC system authentication.
 - NFS Server supports a new buffer alert parameter for the BUFHIGH site attribute. You can use the new parameter to specify a percentage of the storage limit for allowed data buffer space, use the system default, or disable monitoring, and change the specification using the MODIFY operator command. If monitoring is enabled, NFS is designed to issue messages when buffer utilization exceeds the alert threshold, as it approaches the maximum allowed, and when it falls below the threshold. This function is intended to allow you to automate the response to impending buffer shortage events, and help improve NFS Server availability.
 - NFS Server now supports displaying server side statistics for RPC calls using new options on the nfsstat command. This command can be used to display statistics for all NFS servers on a z/OS system, and can make it easier to tune and debug server/client interactions.
 - Statistics about the NFS Client activities to the z/OS NFS Server are displayed. The displayed statistics will include the number of NFS V2, NFS V3, and NFS V4 operations being used by the NFS clients in their communication with the z/OS NFS Server.
- DFSMS enhancements are designed for storage group management and volume selection performance. As volume sizes increase, one percent of a volume represents an increasingly large amount of storage. For example, on a 223 GB volume, 1% is over 2 GB of storage. In z/OS V1.12, the limit on the high threshold you can specify for space utilization for pool storage groups is increased from 99% to 100%. In most cases, IBM recommends a high threshold value less than 100% for storage groups. This allows data sets to expand without an

- increased risk of encountering out-of-space abends. The 100% specification is intended to be used to make more storage capacity available for storage groups that hold static data. Also, SMS processing of volume lists is to be changed in a way intended to improve allocation performance for large volume lists.
- The Integrated Storage Management Facility (ISMF) includes a Data Collection application, DCOLLECT, which provides storage-related measurement data that can be used as input to the DFSMSrmm Report Generator to create customized reports or to feed other applications such as billing applications. In z/OS V1.12, DCOLLECT data class (DC) records are updated to include information about all data class attributes. Also, data set (D) records now include job names, and storage group (SG) records now include information about OAM Protect Retention and Protect Deletion settings.
 - z/OS Communications Server improvements include:
 - The ability to learn indirect prefix routes from IPv6 Router Advertisement messages
 - The ability to associate preference values with all routes that are learned from IPv6 Router Advertisement messages
 - Use of these functions is expected to reduce the number of IPv6 static routes that must be defined and improve the ability to route around network failures when not using OMROUTE to install routes learned via a dynamic routing protocol, such as OSPF.
 - RMF includes information in the CPU Activity Report about how many units of work are running or waiting for a processor (CP, zIIP, or zAAP). Additionally, RMF provides this information in SMF Type 70 records. This new information is expected to be helpful for determining how much latent demand there is for processor time.
 - New RMF Postprocessor reports are now available in XML format. When you specify a new DD name for the Postprocessor output, the Postprocessor will generate the Workload Activity and Device Activity reports in XML format.
 - The Distributor Data Server (DDS) API, which provides access to historical Postprocessor data, is enhanced to grant access to a selection of long-term historical Postprocessor data. Application programs can use URL-based requests for historical RMF data to retrieve information from the DDS. As with RMF Monitor III performance data, the requested Postprocessor data can be provided to requesting programs in a documented XML format for further processing.
 - In z/OS V1.12, Capacity Provisioning Manager (CPM) is enhanced to exploit command correlation. In prior releases, successful or failing command completion events could not be associated with a specific capacity provisioning request. Using command correlation support provided by the hardware is intended to reduce operator intervention when provisioning commands fail for some reason and increase the reliability of hardware management. This function is also available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA30496.
 - In z/OS V1.12 Communications Server, APPN topology database update (TDU) processing is enhanced to provide serviceability enhancements to aid in the

- identification of the network nodes involved in a TDU war, which is the endless exchange of TDUs in contention over the same topology resource, resulting in continuous performance degradation of the APPN network.
- In z/OS V1.12 Communications Server, the DISPLAY TCPIP,,STOR command display and the NMI storage statistics report are enhanced to distinguish the common storage that is used by dynamic LPA for load modules from the ECSA storage that is used for control blocks.
 - In z/OS V1.12 Communications Server, the TN3270E Telnet server is enhanced to:
 - Specify the jobname of the Telnet server issuing a Telnet message
 - Automatically shut down when an OMVS,SHUTDOWN command is issued
 - Pass the connection type (basic or secure) to the application on the CINIT using flags in the CV64 control vector.
 - In z/OS V1.12 Communications Server, VTAM provides applications the ability to pass IP information. An application, such as a session manager, can use this function to inform VTAM and its session partner of any IP characteristics (such as IP address or port number) that are associated with the resource that the application is representing. This enables VTAM displays of the IP information.
 - z/OS Management Facility V1.12 (5655-S28) supports a new RMF-based workload reporting task that provides basic monitoring functions. This systems management task, which requires the optional RMF feature of z/OS, uses RMF distributed data servers (DDSs) running on any supported z/OS release to gather data to display, including most of the data available in RMF Monitor III. You can use it to display information about the systems in one or more sysplexes, including information about connectivity, performance index (PI) status, and multiple measurements, shown in common views each z/OSMF user can save as a performance desktop. A set of predefined desktops is included. Also, this application is designed to gather and display data obtained from an as-is RMF Linux data gatherer tool (rmfpms). rmfpms is intended to gather performance data from Linux on System z and Linux running on Intel-based servers, and is currently available for download. For more information about this tool, see <http://www.ibm.com/systems/z/os/zos/features/rmf/product/rmfhtmls/pmw eb/index.html>
 - You can add a link for the Omegamon XE to the z/OSMF Navigation under the Performance category. In the event of an issue, perhaps identified from the Monitoring Desktops, you could launch Tivoli OMEGAMON XE from the z/OS Management Facility navigator so your system programmers can analyze and diagnose performance problems even faster.

Slide 45

Self-Managing Capabilities 

Improved WLM batch management

- Current and projected resource group maximums considered when starting and rebalancing initiators
- Designed to avoid capping batch address spaces unnecessarily


Workload Management Update for z/OS V1.11 and V1.12
Monday 4:30

New TIMESLICE parameter in IEAOPTxx

- Can help with CPU-intensive discretionary work by letting it get longer time slices when a CPU is available
- Will still be preempted by higher-priority work


IMBED/REPLICATE removal on Recall/Restore

- New function in DFSMSdss and DFSMSshm will remove these attributes automatically during logical RESTORE and recall processing for logically dumped and migrated data sets
- Does not apply to physical dumps:
 - Full-volume dumps
 - Dumps by track
 - Physical data set dumps



45 ©2010 IBM Corporation

Slide 46


Self-Managing Capabilities 

New SMF30 fields to record initiator time

- Initiator address spaces consume processor time on behalf of starting and ending job steps
- Not previously reported time that could be related to a particular job
- Intended to help you better understand the resources consumed by batch jobs and improve the accuracy of chargeback

New Auto-Reply function for WTORs

- Starts very early in IPL
- Specify time limit and default response
- Enable/disable via parmlib or command
- New AUTORxx parmlib members
- Default AUTOR00 member




SMS improvements:

- As volume sizes increase, 1% of a volume is increasingly more storage
 - For example, 1% of a 223 GB volume is more than 2 GB
 - Limit on high threshold for space utilization for pool storage groups increased from 99% to 100%
 - 100% intended to be used for storage groups that hold static data
 - In most cases, IBM recommends a high threshold value less than 100% to allow data sets to expand
- Volume list processing designed to improve allocation performance for large volume lists

46 ©2010 IBM Corporation


Slide 47

Self-Managing Capabilities



DCOLLECT


- Data class (DC) records now include information about all data class attributes
- Data set (D) records now include job names
- Storage group (SG) records now include information about OAM Protect Retention and Protect Deletion settings



47 ©2010 IBM Corporation


Slide 48

Self-Managing Capabilities




DFSMSdss and DFSMShsm now exploit Fast Reverse Restore

- With an IBM System Storage™ DS8000 Series feature
- Designed to allow recovery to be performed from active, original FlashCopy target volumes to original source volumes
- No wait for background copies with full-volume FlashCopy relationships
- DFSMSdss support via a new keyword
- DFSMShsm FlashCopy backup and recovery operations designed to create full-volume FlashCopy relationships for supported devices
- Fast Reverse Restore function intended to support recovery of volumes associated with copy pool backups
 - Includes Space Efficient and Incremental FlashCopy targets
 - DFSMShsm SETSYS parameter allows you to specify whether extent or full-volume FlashCopy relationships are to be established when DFSMShsm invokes DFSMSdss for fast replication backup and recovery



48 ©2010 IBM Corporation

Slide 49


Self-Managing Capabilities 

New callable NMI requests for TCP/IP


- Provide stack network interface information and network interface and global statistics
- Network management applications can use output to monitor interface status and TCP/IP stack activity
- Now supports these new requests:
 - GetGlobalStats - Provides TCP/IP stack global counters for IP, ICMP, TCP, and UDP processing
 - GetIfs - Provides TCP/IP network interface attributes and IP addresses
 - GetIfStats - Provides TCP/IP network interface counters
 - GetIfStatsExtended - Provides data link control (DLC) network interface counters

New SMF119 subtypes for Communications Server


- 032 - DVIPA status change
- 033 - DVIPA removed
- 034 - DVIPA target added
- 035 - DVIPA target removed
- 036 - DVIPA target server started
- 037 - DVIPA target server ended
- 048 - CSSMTP Configuration data records
- 049 - CSSMTP Target server connection records
- 050 - CSSMTP Mail records
- 051 - CSSMTP Spool records
- 052 - CSSMTP Statistics records



Available in SMF data, SMF exit, and NMI (SYSTCPSM)

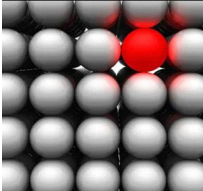
49  ©2010 IBM Corporation

Slide 50

Networking 

Continued focus on IPv6

- IPv4 address exhaustion expected in the 2011-2012 timeframe
- More than you ever wanted to know at: <http://www.potaroo.net/tools/ipv4/index.html>
- Support for security-related RFCs
 - RFC3484, RFC5014, and RFC5175
- Ability to Send DNS Queries Over IPv6
- DFSMSrmm support for IPv6



**IPv6 on z/OS
Tuesday 11:00**

TN3270

- Virtual Storage Constraint Relief: ECSA utilization cut by over 50% since z/OS R7
- Automatic shutdown on F OMVS,RESTART and F OMVS,SHUTDOWN

Trusted TCP/IP Connections

- Within a sysplex with a shared RACF database


Inbound Workload Queueing (IWQ)

- Separate input queue processing for interactive, bulk, and Sysplex Distributor traffic

**What's New in z/OS Communications Server?
Monday 3:00**

**z/OS V1R12 Communication Server - A Detailed Look
Wednesday 3:00**

**Understanding z/OS Communications Server Storage Usage
Thursday 3:00**

50  ©2010 IBM Corporation

Networking

Where would we be without computer networks? Explosive growth in Web-based services, applications, appliances, and mobile devices is fueling a need for increased network performance, scalability, security, and management capabilities.

The z/OS Communications Server is there to meet the challenges with a wide array of networking technologies supported (including both TCP/IP and SNA). System and data security technologies, fault tolerance, autodetection and autorecovery capabilities -- all mean the z/OS Communications Server can provide reliable and trustworthy networking services. With intelligent configuration, dynamic optimization, self-tuning, and network routing, it adapts to different networking conditions and is capable of shifting workloads and traffic to meet quality of service and business needs. Designed for the largest enterprises in the world, z/OS provides network scalability, supporting both IPv4 and IPv6.

In a cluster, the z/OS Communications Server supports multiple applications, tools, databases, operating system images, partitions, servers, locations, and remote locations; can support multiple TCP/IP stacks, and provide different security and networking characteristics for them; can automatically fail over a network; supports dynamic management of networking traffic and can route it by security characteristics, workload priority, and other quality of service characteristics; and applies TCP/IP security capabilities centrally from an attractive, easy-to-use graphic user interface (the Configuration Assistant for z/OS Communications Server in the z/OS Management Facility).

Many data security breaches arise from data being plucked from an unsecured network connection. Internet Protocol Security (IPSec) is just one of the industry standards useful for encrypting packets of a data stream. The z/OS Communications Server already allows for simplified and centralized configuration of IPSec security through its Configuration Assistant and allows most IPSec encryption and decryption to be eligible for the zIIP specialty engine. IPSec encryption on z/OS has the value of encrypting data right at the source. z/OS V1.12 supports Internet Key Exchange version 2 (IKEv2), a more streamlined and efficient method of IPSec dynamic key exchange than the previously available IKEv1. Also for z/OS V1.12, z/OS Communications Server IPSec and IKE support leverage z/OS cryptographic modules that are designed to address the Federal Information Processing Standard (FIPS) FIPS 140-2 security requirements for cryptographic modules. Additionally, z/OS Communications Server IPSec and IKE now support a variety of new cryptographic algorithms, enhanced X.509 digital certificate support, and more. Details on the latest on IPSec and IKEv2 can be found in the Security section. z/OS V1.12 provides many networking performance improvements as well.

Networking performance for interactive workloads can be improved by using new OSA-Express3 Inbound Workload Queueing (IWQ, announced with the new IBM zEnterprise System 196). It is 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). And a streamlined Communications Server execution path for Sysplex Distributor over IWQ is expected to improve performance for sysplex-distributed traffic as well.

Details on the networking improvements for z/OS V1.12 include:

- z/OS Communications Server provides a new DISPLAY TCPIP,,OSAINFO command that you can use to retrieve information about an interface from an OSA-Express feature that is in QDIO mode. The new command is an alternative to using OSA/SF, which lacks information about many of the latest enhancements to the OSA-Express feature and to z/OS Communications Server. This function is limited to OSA-Express3 Ethernet features that are in QDIO mode (CHPID type OSD) and that are running on an IBM System z10 and later System z servers.
- z/OS V1.12 Communications Server is designed to provide notification to the operator console when a Domain Name System (DNS) name server does not respond to a certain percentage of resolver queries during a sliding 5-minute interval. In addition, statistics regarding the number of queries attempted and the number of queries which received no response are displayed for each currently unresponsive name server. This can alert you to a potential problem with your DNS name server configuration that might be adversely affecting applications on your z/OS system. Also, the default value for the TCPIP.DATA RESOLVERTIMEOUT configuration statement, which controls the timeout value for UDP requests sent to a name server, has been changed from 30 seconds to 5 seconds.
- z/OS Communications Server extends the VARY TCPIP,,DROP command to allow the dropping of all established TCP connections for servers that match the specified filter parameters. When issued, each server that is found to match the specified filter parameters will have all its established TCP connections dropped. You can filter by port, jobname, or server ASID. This function is expected to make it easier to move workload from one application instance to another.
- z/OS Communications Server provides the option of keeping a TCP/IP stack isolated from the sysplex; you can use a new configuration parameter to prevent a stack from automatically joining the sysplex group at startup. You can have the stack join the sysplex group at a later time by issuing the VARY TCPIP,,SYSPLEX,,JOINGROUP command.
- z/OS Communications Server is designed to enhance the performance of fast local sockets for TCP connections. This function is automatically enabled.
- z/OS Communications Server provides local and path MTU discovery to learn the correct MTU size for Enterprise Extender (EE) connections. The MTU size is used to modify the link size for EE connections. In z/OS V1.12, z/OS Communications Server is designed to update the link size at the RTP pipe endpoints in addition to the EE endpoints when the MTU size changes.
- z/OS Communications Server packet trace filtering is enhanced to support:
- Including the next hop IP address on the trace output. This can be obtained from the fully formatted packet trace using IPCS. The next hop IP address is also available to applications that consume the real-time packet trace through the real-time TCP/IP networking monitoring API.
- Making packet trace filtering available to encapsulated packets that are used in VIPAROUTE traffic.

- z/OS Communications Server provides the option to check the health of an Enterprise Extender (EE) connection during the activation of the connection. The health of active connections can also be verified.
- z/OS Communications Server is designed to reduce CPU utilization for the TCP/IP Callable Network Management Interface (NMI), EZBNMIFR, GetConnectionDetail. All the filters that are specified for the request must contain the complete identification (4-tuple) of established TCP connections. The 4-tuple of a TCP connection consists of the local IP address, local port, remote IP address, and remote port for the connection.
- The z/OS Communications Server Netstat function provides support for verifying that message catalogs being used are at the correct level when they are opened. This function is intended to prevent Netstat from abending or not functioning correctly when the message catalog is out of synch with the Netstat program.
- z/OS Communications Server enhances TCP/IP data tracing (DATTRACE) to provide two new trace records:
 - A Start record with State field "API Data Flow Starts" that indicates the first data sent or received by the application for the associated TCP or UDP socket
 - An End record with State field "API Data Flow Ends" that indicates the socket has been closed
- z/OS Communications Server supports RFC 3484 by providing a configurable policy table for default address selection for IPv6. The source address selection algorithm and destination address selection algorithm are enhanced to support additional address selection rules in conjunction with the configured or default policy table. For example, you might choose to prefer IPv4 communication over IPv6 by providing a custom policy table for default address selection.
- z/OS Communications Server supports RFC 5014 by providing IPv6 socket API for source address selection. Applications can indicate they prefer temporary IPv6 addresses over public IPv6 addresses or public IPv6 addresses over temporary IPv6 addresses. Additionally, z/OS Communications Server is designed to enhance the SRCIP configuration to allow an administrator to indicate that the TCP/IP stack should prefer public IPv6 addresses over temporary IPv6 addresses. This will allow you to override the preferences specified by an application using the IPv6 socket API for source address selection.
- z/OS Communications Server is designed to allow the system resolver to send requests to Domain Name System (DNS) name servers using IPv6 communication. You can specify IPv6 addresses using the existing NSINTERADDR and NAMESERVER resolver configuration statements in the TCPIP.DATA dataset to define the IPv6 address of the name server.
- In z/OS 1.12, DFSMSrmm supports IPv6.
- z/OS Communications Server allows the coding of MULTIPATH in the TCP/IP profile to enable multipath support for IP packets. You might want this behavior for TCP connections but not for Enterprise Extender (EE) connections. In z/OS Communications Server V1.12, the multipath function is disabled by default for EE connections regardless of the value specified in the TCP/IP profile. You can use the VTAM start option MULTPATH to control the multipath function for EE.

- z/OS Communications Server enhances the digital certificate access server (DCAS) to allow modification of the debug level without restarting the application.

Slide 51

Networking

IKEv2

- IPv4 and IPv6 support
- FIPS-140 mode
- NSSD X.509 support for:
 - Certificate-based signature creation and validation
 - Certificates that contain ECDSA keys
 - Certificate trust chain support
 - IKEv1 or IKEv2 digital signature creation and verification
- New KeyId identity type
- Authentication using pre-shared keys or digital certificates with RSA or ECC keys
- Rekeying and reauthentication of IKE SAs and child SAs
- Hash and URL specification of certificates and bundles
- A new certbundle command to create RFC4306-style bundles

IPSec

- CRL support
- FIPS 140-2 mode

Other improvements for:

- AT-TLS
- Resolver
- EE MTU autotuning
- NSS

Excuse me, do I know you? Certificate Authentication in the Internet Key Exchange

Thursday 1:30

51
©2010 IBM Corporation

Slide 52

Networking

TN3270 server ECSA usage improvement up to and including z/OS V1R12 Communications Server

Release	ECSA for 256K TN3270 sessions
V1R7	798M
V1R8	708M
V1R9	480M
V1R10	440M
V1R11	347M
V1R12 (1)	249M

ECSA for 256K TN3270 sessions

Release	Usage (Megabytes)
V1R7	798
V1R8	708
V1R9	480
V1R10	440
V1R11	347
V1R12	249

The numbers are configuration dependent, but they should give you an idea of the magnitude of the savings achieved in the recent releases.

Note (1): The V1R12 number is a preliminary number based on use of shared ACBs - it may change before general availability of z/OS V1R12 Communications Server

52
©2010 IBM Corporation

Slide 53

Networking

Trusted TCP/IP Connections

- Allow endpoints within a Sysplex or Subplex to establish a trust relationship
 - Exchanges security information that identifies the security context of the other endpoint
 - Without the overhead and CPU-related costs of SSL/TLS with client authentication
 - Requires no application protocol changes
 - Simple API call to the TCP/IP stack
 - Transparent to the client application
 - Security information exchanged using secure XCF messaging
 - Application traffic may take any network path between the client and server
- Current plans are to support these new socket API options for C/C++ (Language Environment), z/OS UNIX System Services Callable (BPXxxxx) and Java
 - Considering whether support will be provided for TCP/IP native socket APIs (EZASMI, Callable, etc.)
 - Implementation is in z/OS V1R12 limited to connections between z/OS systems in a Sysplex within a single RACF security domain

Sharing a global cache – a RACF security domain (a Sysplex)

Application Data Takes Any Path in the Network

53
©2010 IBM Corporation

Slide 54


Networking

Path-length reduction in data processing between AT-TLS and System SSL

- Expected to provide improved performance in encrypted data handling and improve price/performance for workloads using System SSL encryption
- Network CPU overhead for Application Transparent — Transport Layer 3 Security (AT-TLS) can be improved by 30%*.
 - *Based on internal IBM laboratory measurements. The AT-TLS CPU consumption results were obtained on System z10, model 2097-E64. Actual AT-TLS CPU consumption improvement will depend on amount of data being transmitted and whether the workload is interactive or streaming.

54
©2010 IBM Corporation

Slide 55



Networking


Inbound Workload Queueing

- 30-50% increase in network throughput expected for interactive workloads with z/OS V1.12 and OSA-Express3 in QDIO mode with IWQ on z196 and System z10 servers*
- IWQ creates multiple input queues
 - 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, IWQ allows z/OS to use appropriate processing resources for each input queue.
 - Concurrent z/OS processing threads can process each queue, avoiding resource contention.
- In a heavily mixed workload environment, this "off the wire" network traffic separation provided by IWQ reduces processing required to identify and separate work
- Bulk-data (streaming) workloads can also benefit due to fewer network retransmissions (by reducing the incidence of out-of-order packets).
- A streamlined z/OS Communications Server execution path for Sysplex Distributor over IWQ is expected to improve performance for sysplex-distributed traffic

* The interactive networking throughput measurements were obtained on System z10, model 2097-E64. Actual benefit will depend 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. Performance for interactive workloads depends 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.

55 ©2010 IBM Corporation

Slide 56



Networking

Resolver Improvements

- Resolver to monitor DNS server responsiveness
 - When too many requests (default 25%) are not responded to within a 5-minute window, a message will be issued
- Support for IPv6 connections to DNS name servers

Stack Isolation


- New the option to keep a TCP/IP stack isolated from the sysplex
- New configuration parameter
- Can have the stack join the sysplex group at a later time using VARY TCPIP,,SYSPLEX,JOINGROUP

Enterprise Extender maximum transmission unit autotuning

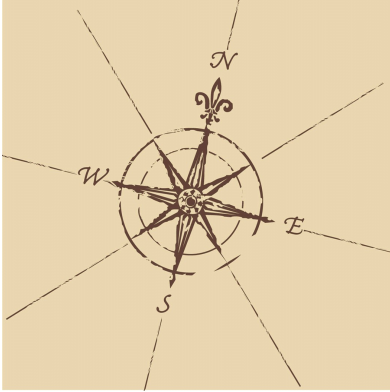
- Local and path MTU discovery to determine best MTU size for Enterprise Extender (EE) connections
- MTU size used to modify link size for EE connections
- In z/OS V1.12, the link size will be updated at the RTP pipe endpoints in addition to the EE endpoints when the MTU size changes.

56 ©2010 IBM Corporation

Slide 57




Statements of Direction



57

©2010 IBM Corporation

Slide 58



Statements of Direction*

Reminders*:

- IBM plans to discontinue software delivery on 3480, 3480 Compressed (3480C), and 3490E tape media.
- Concurrently with z/OS V1.12 availability, ServerPac and CBPDO; the SystemPac®, ProductPac®, and FunctionPac fee-based offerings; and selective follow-on Service (SFS) for the fee offerings will be orderable on DVD.
- In the first quarter of 2011, ordering the COD on DVD will be supported.
 - IBM recommends using Internet delivery, but DVD support may provide an option for those who cannot use it
- IBM plans to remove the Enhanced PSP Tool, host compare program, and the associated extract files from the IBM Technical Support Web site (software.ibm.com/webapp/set2/psp/srchBroker) at the end of 2010.
 - Function replaced by FIXCAT HOLDDATA and SMP/E's REPORT MISSINGFIX command
 - SMP/E function available for all supported releases; Enhanced HOLDDATA is available in all IBM service delivery offerings and in the 730-day file from the Enhanced HOLDDATA Web site at: service.software.ibm.com/holddata/390holddata.html


58

* 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

Slide 59

Statements of Direction*



Reminders:

- The native SSL/TLS support in the TN3270 server and FTP client and server will not be enhanced to address FIPS 140-2 requirements
 - Use AT-TLS instead
- IBM intends to withdraw BIND DNS 9.2.0 function from z/OS Communications Server
 - IBM recommends you use the Resolver cache function in z/OS V1.11 to cache DNS responses
 - If you use the z/OS BIND 9.2.0 function as a primary or secondary authoritative name server, you should investigate using BIND on Linux for System z


In a future release, IBM plans to withdraw support for zFS multi-file system aggregates

- When this support is withdrawn, only zFS compatibility mode aggregates will be supported. (A zFS compatibility mode aggregate has a single file system per data set.)

69 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

Slide 60

Statements of Direction*




Reminders:

- Infrastructure for RTLS for Language Environment
 - Last release to support Run-Time Library Services z/OS V1.5
 - In z/OS V1.12, IBM removed the underlying CSVRTLS services
 - A way to track usage so you can find any programs that using these services available for z/OS V1.9, V1.10 with the PTF for APAR OA29019
- msys for Setup element of z/OS has been removed from z/OS V1.12
 - IBM intends to continue to deliver improvements to help with z/OS setup and configuration in the future
- IBM no longer plans to withdraw support for VSAM IMBED, REPLICATE, and KEYRANGE attributes in a future release
 - Based on customer feedback
 - IBM recommends that you stop using these attributes
 - IBM still plans to remove IMBED and REPLICATE attributes during logical DFSMSdss restore operations and DFSMSshm recall operations as announced in IBM United States Software Announcement 207-175, dated August 7, 2007

60 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

Slide 61

Statements of Direction*




Reminders:

- IBM plans to pursue an evaluation to the FIPS 140-2 using NIST's Cryptographic Module Validation Program (CMVP) for the PKCS #11 capabilities of ICSF. The scope of this evaluation will include algorithms provided by the CPACF that are utilized by ICSF. This is intended to help satisfy the need for FIPS 140-2 validated cryptographic functions when using z/OS Communications Server capabilities such as the IPSec protocol.
- Plans related to Extended Address Volume (EAV) larger volume sizes as described in 5694-A01, Preview: z/OS V1.10, announced in Software Announcement 208-042, dated February 26, 2008, will be communicated at a later date.

61 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

Slide 62

Statements of Direction*




New news

- IBM plans to add support for Microsoft® Windows® 7 64-bit Professional Edition to z/OS Management Facility V1.12 with the 32-bit version of Internet Explorer 8 and Mozilla Firefox 3.5. IBM plans to add support for the 32- and 64-bit versions of Windows 7 Professional Edition to DFS SMB Client, Hardware Configuration Manager, and NFS.
- z/OS V1.12 is the last release to include z/OS Distributed Computing Environment (DCE) and Distributed Computing Environment Security Server (DCE Security Server).
 - IBM recommends the IBM WebSphere® Application Server, the IBM Network Authentication Service, and the IBM Directory Server as replacements. See the DCE Replacement Strategies Redbook for more details:
<http://www.redbooks.ibm.com/redbooks/pdfs/sg246935.pdf>

62 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

Slide 63


Statements of Direction* 


New news

- z/OS V1.12 is the last release to include z/OS Distributed File Service support using the Distributed Computing Environment (DCE) architecture
 - IBM recommends the z/OS Network File System (NFS) as a replacement
 - Support in Distributed File Service for Server Message Block (SMB) architecture remains, and is not affected by this withdrawal of support
- In a future release, IBM plans to remove the capability to change the default Language Environment run-time options settings via SMP/E installable USERMODs. IBM recommends using the CEEPRMxx parmlib member to change the default Language Environment run-time options for the system.
 - Support for NONOVR added in z/OS R12
- z/OS V1.12 is planned to be the last release to include the z/OS UNIX System Services Connection Scaling functions provided by the Connection Manager and Process Manager components.


63 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

Slide 64






The Future Runs on System z



Optimize your z/OS environment



64