Slide 1

IBM Systems & Technology Group

What's New in z/OS?

Bob Rogers
John Eells
SHARE 107, Session 2838
August 14, 2006

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.

ON DEMAND BUSINESS™

© 2004, 2006 IBM Corporation

Slide 2

Trademarks

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

AIX*	ES/9000*	Language Environment*	S/390*
AnyNet*	FICON*	Lotus*	Sysplex Timer*
CICS*	FlashCopy*	Multiprise*	System Storage
DB2*	GDPS*	Notes*	System z
DB2 Connect	Geographically Dispersed Parallel Sysplex	OMEGAMON*	System z9
DB2 Universal Database	HiperSockets	On demand business logo	SystemPac*
developerWorks*	Hiperspace	OS/390*	Tivoli*
DFSMSdfp	HyperSwap	Parallel Sysplex*	TotalStorage*
DFSMSdss	IBM*	PR/SM	Virtualization Engine
DFSMShsm	IBM eServer	Processor Resource/ Systems Manager	VTAM*
DFSMSrmm	IBM e(logo)server*	pSeries*	WebSphere*
DFSORT	IBM logo*	RACF*	z/Architecture
Domino*	IMS	RAMAC*	z/OS*
DRDA*	Infoprint*	Redbook	z/VM*
Enterprise Storage Server*	IP PrintWay	RMF	zSeries*
ESCON*	iSeries		

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of the Intel Corporation in the United States and other countries.
Linux is a trademark of Linux Torvalds 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.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

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

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

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

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

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

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

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.

Slide 3

Agenda

- **z/OS® Business Overview – Bob Rogers**
- **z/OS 1.8 Update – John Eells**

Slide 4

Over 93% on supported releases of z/OS.

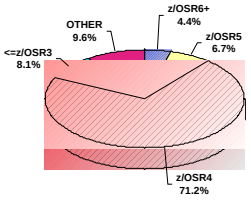

YE 2004

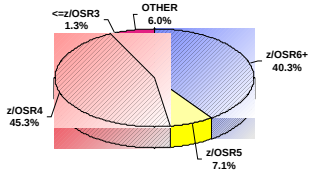

YE 2005

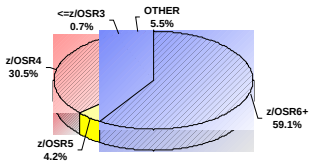

2Q06

Percent (%) of z9 EC and z9xx WW System Inventory

Slide 5

IBM System z9

The server designed to help protect, grow, and meet the demands of enterprises of all sizes

The IBM System z9™ Enterprise Class (z9 EC) – formerly called z9-109 – and the new IBM System z9 Business Class (z9 BC) deliver excellence in enterprise computing and are designed and optimized for on demand business

z9 EC

- **Built on more than 40 years of industry leadership and taking that leadership to new levels**
  - Scalability
  - Availability
  - Security
- **New roles for mainframe**
  - A mainframe for everyone
  - Helping to drive increased value from data and applications including the announcing of the availability of System z9 Integrated Information Processor (zIIP)
  - Helping to simplify management and reduce costs of storage subsystems with new connectivity options

z9 BC

Now there is a System z9 for everyone

Slide 6

System z9 EC . . . built to help protect and grow with your business

- **Capacity to meet your business objectives**
  - Capacity on demand to enable minimal downtime
  - Large mainframe server in a single footprint with the S54
  - Leadership capabilities to help improve I/O access *
- **Helping lower risk of downtime**
  - Leader in Parallel Sysplex® clustering and GDPS® services
  - Superior access to applications via comprehensive protection from unplanned and planned outages
- **Security features to help address compliance**
  - Industry certifications and regulatory compliance
  - Leadership capabilities in cryptography and encryption

Now the System z9 BC also available – but in a smaller package

* Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the numbers stated here.

Slide 7

z9 BC – The modern mainframe for the small to medium enterprise

- **Based on System z9 technology**
- **Designed for flexibility in 2 new models**
- **More engines for more workloads**
  - IBM System z™ Application Assist Processor (zAAP), Integrated Facility for Linux (IFL), Internal Coupling Facility (ICF), zIIP
- **On demand upgrade capability**
  - Exceptional upgradeability
  - On/Off Capacity on Demand (On/Off CoD) functions available
- **Enhanced networking and connectivity options**
- **Built with System z9's cryptographic and encryption functions**
  - ATM/POS Remote Key Load
- **EWLC and Tiered EWLC Software Pricing Structure**
- **Operating system support – similar to z9 EC, plus z/OS.e**

Low entry point and more choices

Slide 8

Technology Evolution with Mainframe Specialty Engines

Building on a strong track record of technology innovation with specialty engines, IBM introduces the System z9 Integrated Information Processor

IBM System z9 Integrated Information Processor (IBM zIIP) 2006

Designed to help improve resource optimization for eligible data workloads within the enterprise

System z9 Application Assist Processor (zAAP) 2004

Incorporation of Java™ into existing mainframe solutions

Integrated Facility for Linux® (IFL) 2001

Support for new workloads and open standards

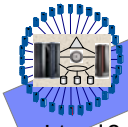

Internal Coupling Facility (ICF) 1997

Centralized data sharing across mainframes

Slide 9

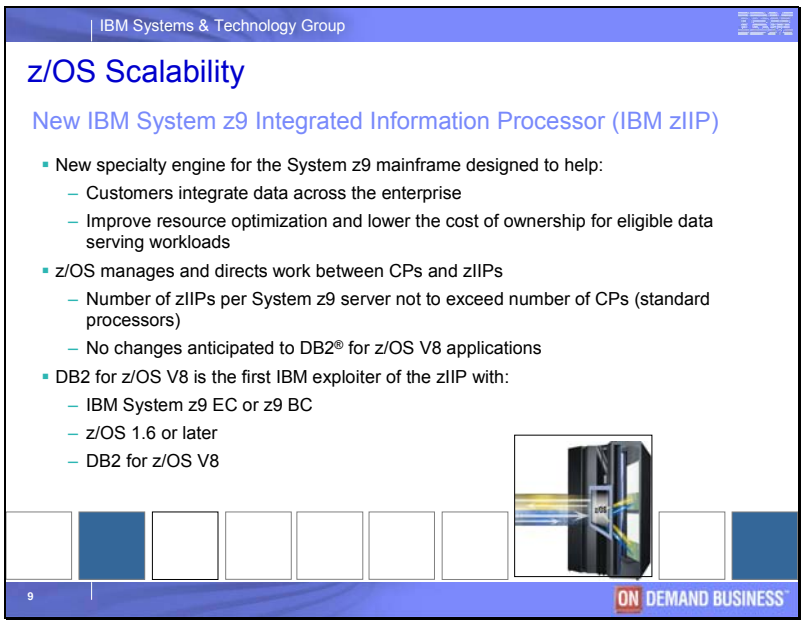

IBM System z9 Integrated Information Processor (IBM zIIP): In a press release on January 24, 2006, IBM previewed a new specialty engine called the IBM System z9 Integrated Information Processor (IBM zIIP) for the z9 EC. The IBM zIIP is the latest customer-inspired specialty engine for the IBM System z9 mainframe. Following on the success of the widely accepted Integrated Facility for Linux (IFL) and System z9 Application Assist Processors (zAAP), the zIIP is designed to help improve resource optimization and lower the cost of eligible workloads, enhancing the role of the System z9 mainframe as the data hub of the enterprise.

When available, the zIIP's execution environment will accept eligible work from z/OS V1.6 or higher. This will manage and direct the work between the z9 general-purpose processor and the zIIP.

The zIIP is designed so that a program can work with z/OS to have all or a portion of its Service Request Block (SRB) dispatched work directed to the zIIP. The z/OS operating system, acting on the direction of the program running in SRB mode, controls the distribution of the program's work between the general-purpose processor and the zIIP.

DB2 for z/OS V8 (5625-DB2) will exploit the zIIP capability for eligible work. Types of eligible DB2 workloads executing in SRB mode, portions of which can be directed to the zIIP, are:
Workloads that leverage existing multi-tiered DB2 applications that via SQL calls use DRDA over a TCP/IP connection to access the DB2 database server from the application
Workloads that offer new opportunities to deploy data warehousing and BI solutions that utilize star schema parallel queries on the System z9 servers
Select internal DB2 for z/OS V8 utility functions (Load, Reorg, Rebuild Index) written in SRB mode, that are used to maintain index maintenance structures

For more information, visit
http://www.ibm.com/systems/z/feature012406/"ibm.com/systems/z/feature012406

Slide 10

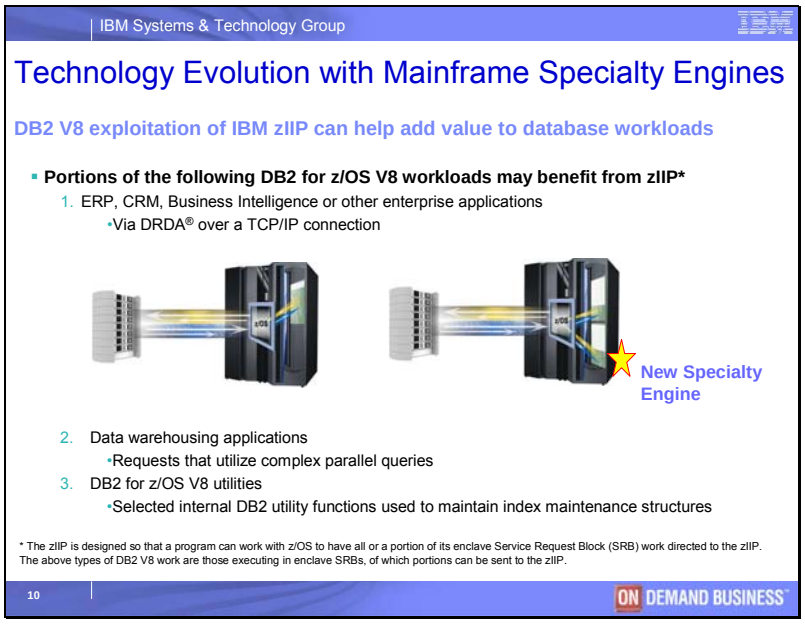

Slide 11

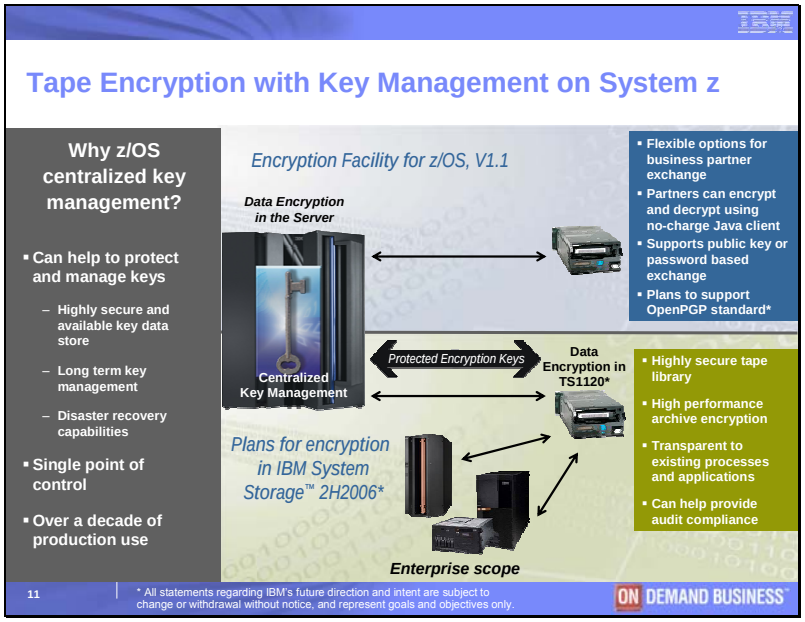

Slide 12

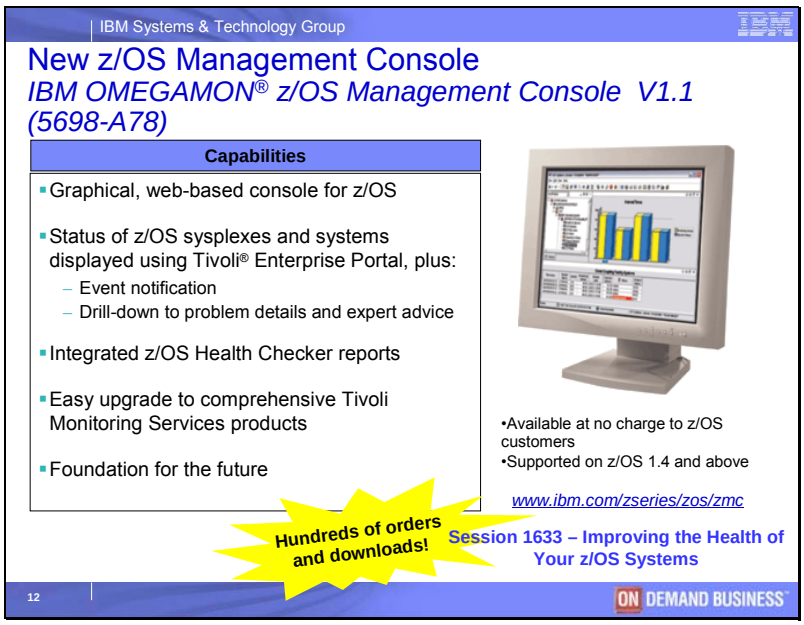

A major part of our simplification of z/OS management is built around a new user interface – the z/OS Management Console.

Information Technology workers by automating, eliminating, and simplifying many z/OS management tasks. The first
phase of the new user interface is designed to provide real-time health check information executed by the IBM Health Checker for z/OS and configuration status information for z/OS systems and sysplex resources. The new interface is intended to form the base for providing built-in automation and expert advice capabilities for z/OS management tasks. It is designed to provide detailed contextual information on alerts and corrective actions.

Slide 13

Slide 14

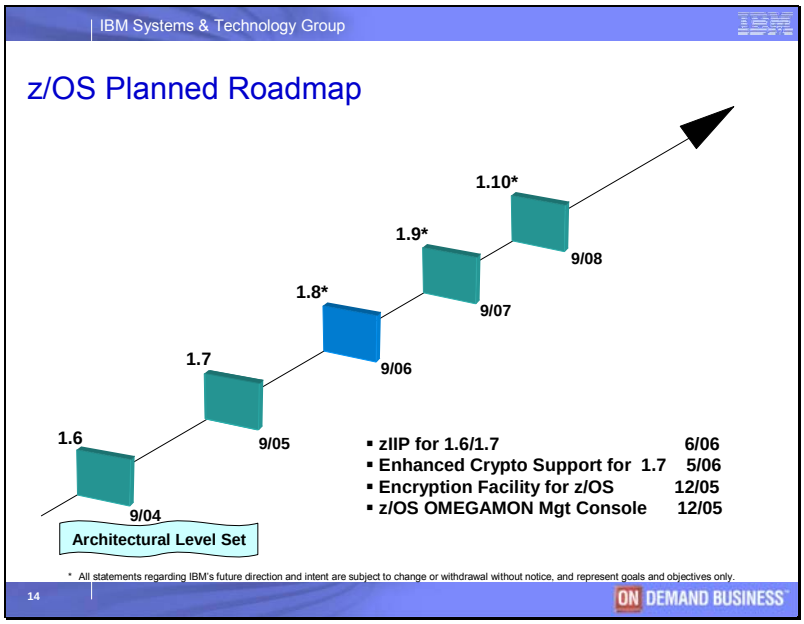

Slide 15

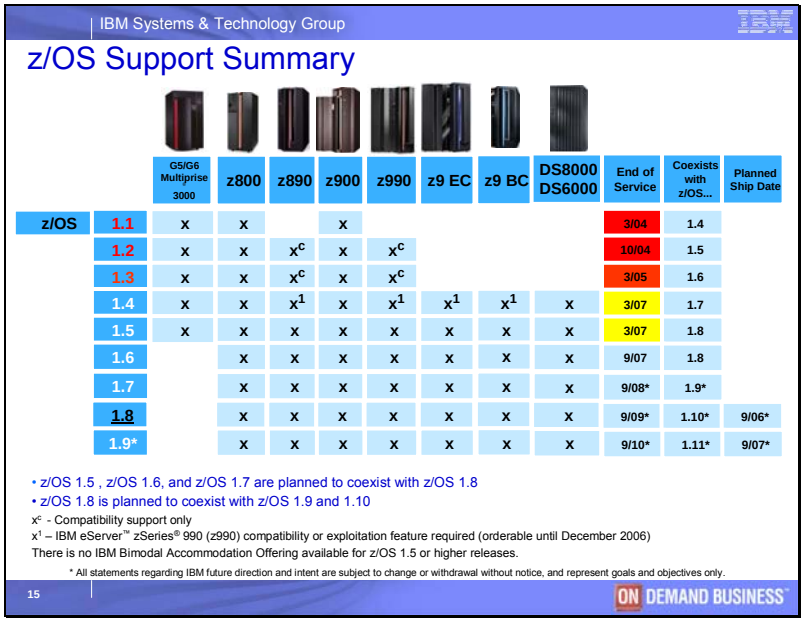

This chart summarizes much of the information about z/OS that you need to know. Server and DASD hardware support, end of service dates, coexistence, and planned availability dates are shown here for existing and planned releases through z/OS R9. Note that:

- The lowest supported level of z/OS is R4
- R8 is the currently orderable release.

Slide 16

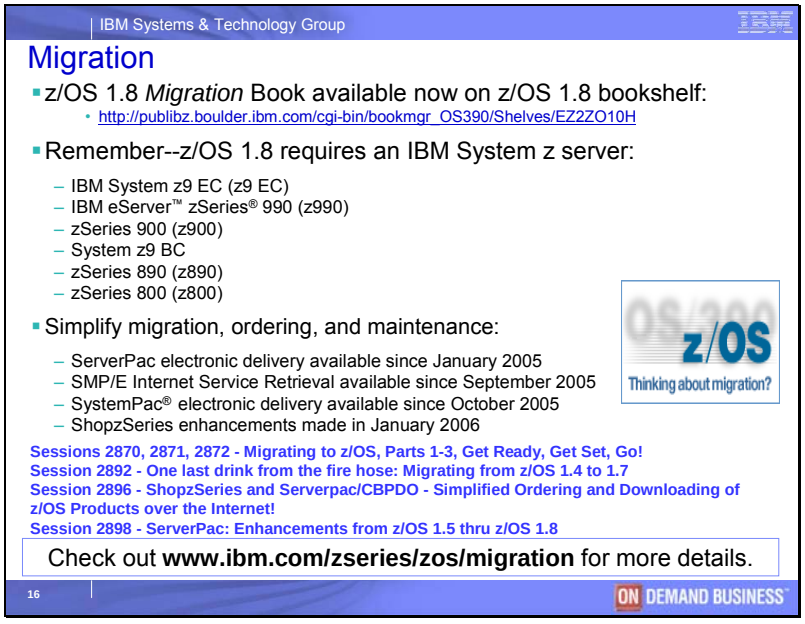

Slide 17

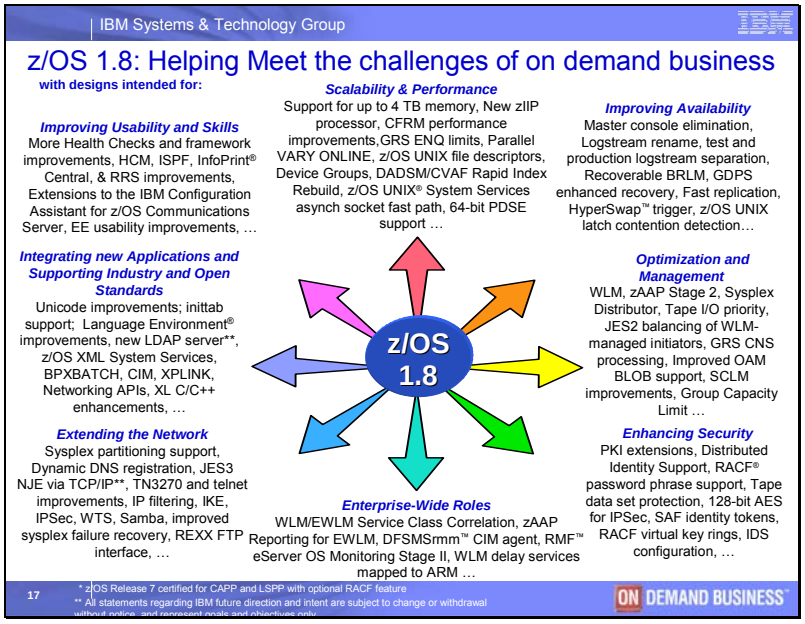

Slide 18

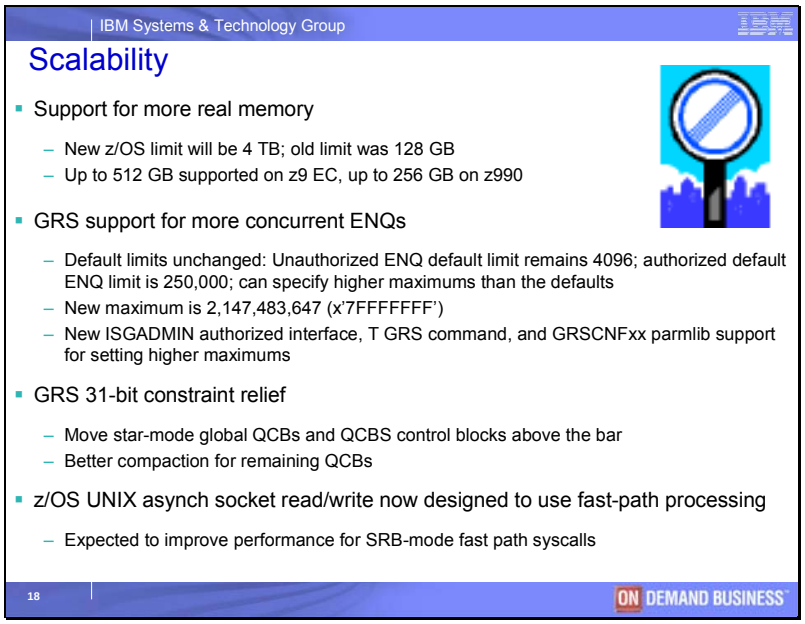

Scalability

z/OS V1.8 helps extend system limits and continue to support application workload growth. There is support for up to 4 TB of real memory on a single z/OS image. This can allow the use of up to 512 GB of real memory on a single z/OS image on IBM System z9 servers and up to 256 GB on z990 servers. There are also improved support capabilities

for large-format sequential data sets, increased system-wide and intra-address space limits for Global Resource Serialization (GRS), and improved high-level language support for VSAM data sets with extended addressability.

New and enhanced z/OS V1.8 function:

There is **support for up to 4 TB (4,398,046,511,104 bytes) of real memory** on a single z/OS image on IBM System z9 servers and up to 256 GB on IBM zSeries z990 servers, an increase from the prior maximum of 128 GB. This can enable programs that use large amounts of real memory to avoid paging and swapping overheads, and helps enable workload growth.

GRS offers a new programming interface designed to allow authorized programs to set their own **concurrent enqueue (ENQ) limits** within their address spaces. Some subsystems, such as DB2, often need to open a large number of data sets concurrently. Others, such as CICS, can manage many different units of work within a single address space. This new support can allow such programs to increase their ENQ limits beyond the system's default limit. Also, this support is designed to allow system programmers to dynamically alter the system-wide limits.

GRS design is changed to **move selected GRS control blocks above the 2-GB bar**. This helps avoid GRS storage constraints as the storage size, processor capacity, and workloads running on large z/OS images grow.

PDSE exploitation of 64-bit addressing. These changes are expected to allow many more PDSE members to be open at the same time.

Language Environment delivers enhancements to the C run-time library functions **fgetpos(), fsetpos(), and fseek()**. This can improve the performance of repositioning operations within multivolume data sets.

Language Environment provides z/OS XL **C/C++ run-time library support for large-format sequential data sets** opened using QSAM (noseek in C/C++). This support removes the constraint of 65535 tracks per volume for sequential data sets.

Language Environment provides z/OS XL **C/C++ run-time library support for VSAM data sets with extended addressability**. This support is for key-sequenced (KSDS), entry-sequenced (ESDS), and relative-record data sets (RRDS).

DFSMSrmm extends support for **managing removable media across the enterprise.** The DFSMSrmm CIM agent is designed to support the creation, change, and deletion of volumes and data sets. This is in addition to the query and display capabilities provided in z/OS V1.7.

The CIM agent uses the **OpenPegasus CIM Server**. It is able to run on z/OS and other systems supported by OpenPegasus. In addition, to improve systems support across the

enterprise, DFSMSrmm support uses common time and provide support for displaying and setting dates and times in any chosen time zone.

Support for **rapid index rebuild** in z/OS V1.8 is designed to help speed VTOC conversions from nonindexed to indexed

VARY command processing is changed to process **VARY ONLINE commands in parallel**, complementing the parallel VARY OFFLINE processing introduced in z/OS V1.7. Parallel VARY processing is designed to decrease elapsed time for commands used to change the status of many devices.

Slide 19

Scalability

- CFRM performance enhancements Stage 2
  - Designed to allow more systems, structures, and connectors to be added without availability impacts
- Parallel VARY ONLINE processing
  - Designed to reduce the duration of VARY commands for large numbers of devices
  - Expected to reduce serialization contention
  - Complements parallel VARY OFFLINE processing, available with z/OS 1.7
- Language Environment support for sequential data sets larger than 64K tracks and for VSAM extended addressability data sets
  - QSAM support in Language Environment for C/C++ programs (using **noseek**)
  - Support for ESDS, KSDS, RRS extended format data sets with extended addressability

Session 8221 – What's new in Language Environment for z/OS

Slide 20

Scalability

- z/OS UNIX limit on file descriptors per process increased from 128K to 512K
  - MAXFILEPROC limit will be 524288 in BPXPRMxx and SET OMVS command
  - Can restrict for individual users by setting FILEPROCMAX in the OMVS segment
  - C and assembler programming interfaces support new limit
- Support for more than 32K device groups
  - Device groups are consumed by tape devices (1 per device) and DASD esoteric names (1 per device range per esoteric name).
  - You can reach this limit with fewer than the 63.75K maximum number of devices
  - APAR OA02983 increased the limit to 64K-1 devices (x'FFFF') at the end of 2003
  - The new limit will be really high (4G-1, or x'FFFFFFFF')
  - Note: You can check on your systems by looking at the number of entries in the Group Pointer Table using the IPCS LISTEDT command.
- DADSM/CVAF Rapid VTOC Index Rebuild
- 64-bit exploitation by PDSE
  - SMSPDSE and SMSPDSE1 will use storage above the bar for control blocks
  - New SMS initialization parameter to specify amount of storage to use above the bar
  - Relieves the prior limitation of about a million open PDSE members
  - Can allow over a million concurrently-open PDSE members
  - Also, option to retain directory and member cache for closed data sets and SMF 14/15 cache statistics

Session 2919 – What's new in z/OS UNIX?

Slide 21

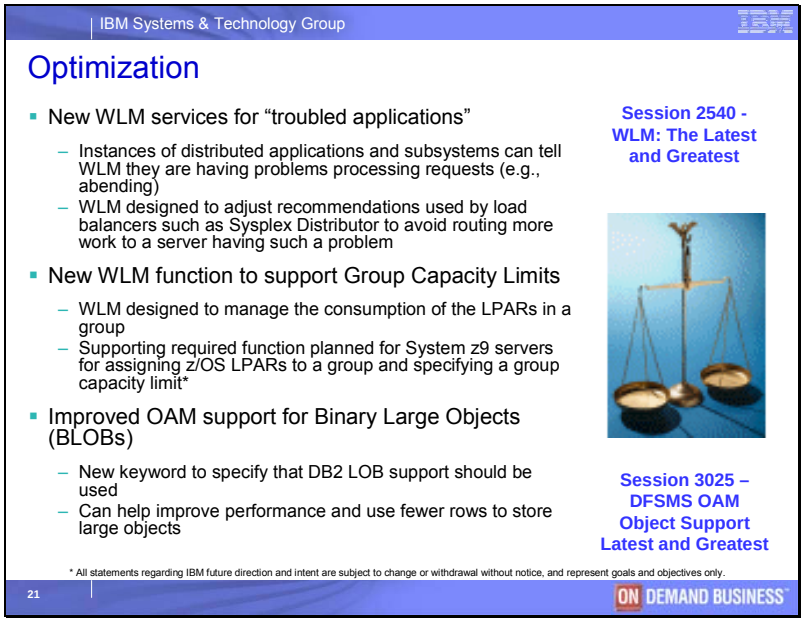

Optimization and management capabilities

In z/OS V1.8, Communications Server offers additional network and workload balancing improvements by favoring local servers when possible to help reduce inter-CEC network traffic. There is also enhanced Workload Manager (WLM) processing for zAAP workloads, improved I/O priority for tape devices, and more integration between WLM and Enterprise Workload Manager (EWLM). These functions are intended to help optimize workloads within an image and across a sysplex:

In z/OS V1.8, Enhancements to Workload Manager services can allow **z/OS applications and middleware to report to WLM abnormal operating conditions that have an impact on their ability to successfully process new work requests.** WLM is designed to use this information to influence the workload balancing recommendations it provides to functions such as Sysplex Distributor and the Load Balancing Advisor. When problem conditions are reported to WLM, workload balancing recommendations are expected to be appropriately reduced, allowing these load balancing functions to avoid routing new work requests to the target applications that are experiencing these conditions.

A new option is designed to allow the z/OS **Sysplex Distributor to favor local system target servers where possible**, while avoiding servers that are no longer active or are overloaded. If the Sysplex Distributor chooses a target on the same system as the client, this new design optimizes the connection in the following manner:

- Traffic for the target connection is no longer routed to the sysplex distributor routing stack.
- The connections are eligible for the "fast local sockets" optimized path.
- Sysplex Sockets report these sockets as being on the same system.

In z/OS V1.8, Workload Manager (WLM) improvements:

- **Manage workloads based on zAAP delay** in addition to CP delay. This is designed to manage workloads based on both zAAP and CP utilization.

- Change the way the **IFAHONORPRIORITY=YES** parameter in IEAOPTxx parmlib members is processed. This function is designed to enable zAAPs to obtain help from CPs when there is more zAAP eligible work than zAAP capacity. The change helps limit the amount of work that is scheduled on CPs when sufficient zAAP capacity is available to run work that is eligible to run on zAAPs. The changed function is also available with APARs OA14131 and OA13953 for z/OS and z/OS.e V1.6; and z/OS and z/OS.e V1.7.

- System Resource Manager (SRM) calculates a **new static I/O priority** for all address spaces and enclaves. The new I/O priority is derived from the importance of the service class the unit of work is classified to, and is designed to be used by IOS when no dynamic I/O priority is provided by WLM or SRM. This function is supported **for tape** I/O only.

Global Resource Serialization (GRS) enhancements **enable GRS Star users to specify which system will be the contention notifying system (CNS)**. In a GRS Star environment, global contention is reported by a single image acting as the CNS. Allowing you to specify the placement of CNS work can help you better balance workload in a sysplex. All systems in a sysplex must support this function to enable its use. In addition to z/OS V1.8, this new function is included on z/OS V1.7 with APAR OA11382.

EWLM support is added to z/OS WLM:

- **WLM is designed to accept the classification of work from EWLM.** EWLM transaction classes can be correlated to WLM service classes via classification rules in the WLM Service Definition. If such a correlation has been defined, WLM assigns the specified WLM service class to the end-to-end work in the EWLM transaction class rather than reclassifying it. A performance administrator on the EWLM control center is then be able to see integrated information about end-to-end work within the same transaction class for reporting and management. This helps improve cross-platform workload management.
- With z/OS V1.8, **a high-performance EWLM instrumentation option via WLM Enclave Services** is introduced that is designed to help reduce the instrumentation overhead significantly. The high-performance EWLM instrumentation option is suitable for applications that implement a sequential transaction processing flow without sub-transactions.
- **WLM execution delay monitoring services are mapped to appropriate Application Response Measurement (ARM) services.** This simplifies the

EWLM instrumentation of subsystems that are already instrumented with WLM execution delay monitoring services, such as CICS and IMS.

WLM has been enhanced to allow you to **specify the capacity of system-wide resource groups** as a percentages of a single processor's capacity or as a percentage of an LPAR's capacity. This function is intended to allow you to specify relative capacities that change dynamically with the capacity of a processor or LPAR.

The function originally planned for the SYSSTC1-SYSSTC5 service classes as described in Software Announcement 206-039, (RFAXXXXX) dated February 28, 2006, is not included in z/OS V1.8. The classes themselves remain, but they are reserved for possible future use.

In the past, only CP data was reported to EWLM. With z/OS V1.8, **zAAP data is included in the CPU using and delay samples** as well as in the CPU service times reported to EWLM for processes using zAAPs.

The Object Access Method (OAM) **DB2 Binary Large Object Support** enables objects larger than 32 KB to be stored using DB2's large object (LOB) support and the binary large object (BLOB) data type. A new keyword is supported on the OAM1 statement of the IEFSSNxx parmlib member to indicate that DB2's LOB support should be used. Storing in LOB storage structures improves performance when working with large objects. A 256-MB object, for example, can be stored in a single row in a LOB storage structure, but would require 8000+ rows in a conventional 32-KB table.

There is coexistence support for earlier release levels to coexist in an OAMplex with z/OS V1.8. Note: OAM at z/OS V1R8 requires a minimum level of DB2 V7 even if OAM DB2 Binary Large Object Support is not exploited.

OAM Object Tape Enhancements adds automatic selection of RECYCLE-eligible tape volumes to the existing MOVEVOL with RECYCLE function, and provides support for an immediate backup copy to be created for an object at the time the object is originally stored.

The **z/OS UNIX System Services asynchronous socket read and write operations** (the aio_read() and aio_write() callable services) are converted to use fast-path processing. In addition, performance is improved for all fastpath syscalls that are called in SRB mode. This improves performance for applications that use asynchronous I/O or that run in SRB mode.

SCLM design changes to reduce the number of ISPF service calls and cache certain data, improve the performance of the SCLM Library utility (option 3.1), are included in z/OS V1.8.

VARY processing is redesigned to bring multiple devices online in parallel. This provides a significant improvement in the elapsed time required to vary a large number of devices online.

JES2 is designed to help balance workload in a multi-access spool configuration within a sysplex. This new design is intended to use approximately the same percentage of active WLM-managed initiators in each service class on each system. Depending on the characteristics of your workloads and their arrival patterns, this can help balance batch workloads across a MAS configuration within a sysplex.

Slide 22

Optimization

- Improved WLM IBM zAAP support
  - WLM designed to manage based on both CP and zAAP delays
- Improved IFAHONORPRIORITY=YES processing
  - New design intended to use CP resource only when zAAP workload goals are expected to be missed (also available with APARs OA14131, OA13953, and OA15297 on z/OS 1.6 and 1.7)
- Sysplex Distributor enhancements
  - Designed to favor local servers when possible
  - Builds on integration with WLM introduced with z/OS 1.7
- JES2 load balancing for WLM-managed initiators
  - Designed to keep approximately the same percentage of WLM-managed initiators in each service class in use on all systems in a sysplex within the same MAS
  - Designed to help balance sysplex batch workloads

Session 3967 – Sysplex Network Technology Update

Session 2655 – JES2 z/OS 1.8 Product Overview & Review of Selected 1.7 Features

Slide 23

Optimization

- I/O priority improvements
  - SRM designed to assign I/O priority for tape when there is no WLM recommendation
- Ability to specify CNS notification system for GRS Star
  - Also planned for z/OS 1.7 with APAR OA11382*
- WLM delays mapped to ARM services

Session 2884 – GRS: Forum at 6

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

Slide 24

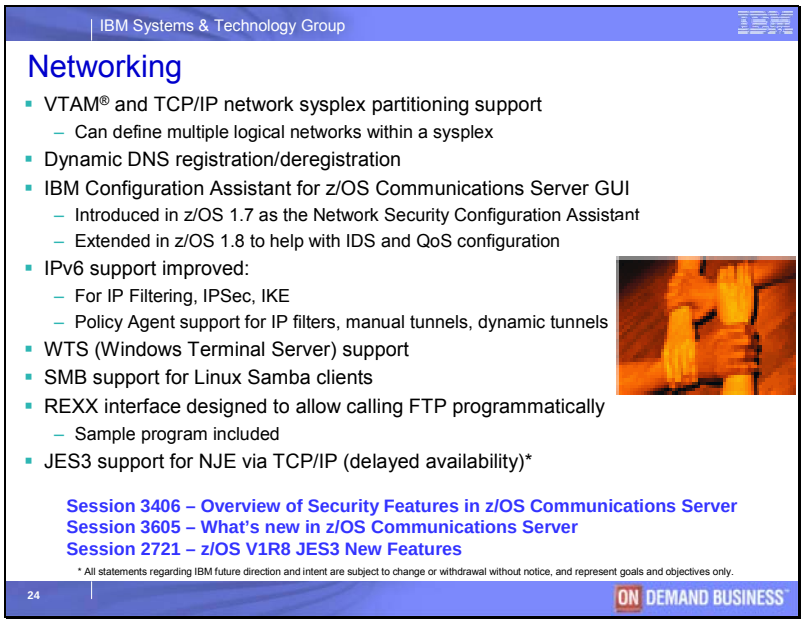

Networking

In z/OS V1.8, there are significant improvements for networking and communications, including the ability to specify subplexes for TCP/IP communications and improved sysplex autonomic functions, support for JES3 NJE connections via TCP/IP, improvements in networking security, and support for Windows Terminal Server (WTS).

The following enhancements are for Sysplex environments:

- The ability to specify a **subdivision of a sysplex into multiple "subplex" scopes** from a sysplex networking function perspective is delivered. For example, some VTAM and TCP/IP instances in a sysplex might belong to one subplex, while other VTAM or TCP/IP instances in the same sysplex might belong to other subplexes. This function can be useful in scenarios where different LPARs within the same sysplex need to be isolated into different security zones, with isolation between those zones.
- The sysplex autonomics function is enhanced to deliver **monitoring of critical network interfaces** so that sysplex autonomics recovery can be triggered when a failure occurs at the interface.
- A new application allows **dynamic registration and deregistration of Domain Name Server (DNS) records** for z/OS hosts, host groups, servers, and server groups and their associated zones based on their current availability. This function provides an improved alternative to the dynamic registration and deregistration of DNS hostnames currently available with the z/OS BIND DNS 4.9.3 server.

Support for **JES3 NJE communications using TCP/IP** is planned to be made available in z/OS V1.8. This function includes support for IPv6, secure sockets (SSL/TLS), and all the NJE constructs (ENDNODE, SUBNET, Store-and-Forward) supported by the owning JES. This new support is in addition to the SNA and BSC protocols currently supported

by JES3. In order to use the NJE/TCP support, both sides of the connection are required to support NJE/TCP. This function is planned to be made available during the first half 2007.

Functions to support **IP filtering, IPSec, and Internet Key Exchange (IKE) for IPv6** are provided:

- The Policy Agent, to help configure IP filters, manual tunnels, and dynamic tunnels for IPv6
- The TCP/IP profile, to allow configuration of default IP filters for IPv6 when the policy-based IP filters are not active
- The z/OS Communications Server IKE daemon, to negotiate dynamic tunnels for IPv6
- The IBM Configuration Assistant for z/OS Communications Server GUI, to help configure IPSec for IPv6 in the Policy Agent and the IKE daemon
- The ipsec command, to display and modify installed IP filter information, manual tunnel information, and dynamic tunnel information for IPv6
- The Traffic Regulation Management daemon (TRMD), to support logging of IP Security events, such as IP filter permits and denies, for IPv6

z/OS Communications server delivers a rich set of **callable APIs** that can be used by network management applications. In z/OS V1.8, a number of additions and improvements are made to this set of APIs:

- There is an API for dropping multiple TCP connections or UDP endpoints.
- The existing Enterprise Extender network management interface is enhanced to allow the specification of wildcard characters in the CP name on HPR connection requests. For applications using this interface to gather HPR connection data, using a wildcard value can reduce the number of NMI requests issued to obtain all of the data, and this improves application performance.
- The ability to manage TN3270 connections is improved with:
- The provision of a new API designed to allow for the retrieval of performance data for TN3270 server sessions
- The addition of the TCP connection ID to the TN3270 SMF records

WTS support: In the past, the z/OS SMB server could handle a single user session over a single communications session from a Windows PC. The WTS is able to act as a client to the z/OS SMB server. The WTS allows many Windows clients to connect to it. The server, in turn, act as a client to the z/OS SMB server and send many user sessions over a single communications session.

Linux Samba client support: The z/OS SMB server also supports Linux Samba clients.

There are the following changes for SNA and Enterprise Extender (EE):

- An **EE connectivity test command** is provided to assist in debugging various network problems. This new test command is intended to be used to test an existing EE connection, or to assist in determining why an EE connection cannot be established.

- A number of miscellaneous usability, serviceability, and problem determination improvements are made to SNA, and in particular EE.

FTP enhancement: z/OS V1.8 Communications Server provides a new REXX interface designed to be used to invoke the FTP client programmatically. This API support extends the existing FTP Client API to support the REXX programming language. A sample REXX program is delivered.

Telnet enhancements:

- Enhancements allow you to specify that a TN3270 server should automatically clean up hung SNA Telnet sessions when a new TN3270 connection is initiated. This is designed to help reduce the number of reconnect failures caused when Telnet still has a SNA session for the original connection.
- There is support for the resolution of system symbols for Telnet unformatted system services message (USSMSG) processing. System-specific symbols such as the system name are displayed to help you diagnose problems.
- The USS table assignment from the LU exit is designed to provide more flexibility in assigning a USS table based on client criteria.

z/OS V1R6 Communications Server and subsequent releases include a **standalone TN3270 Server**. This standalone TN3270 server can provide increased flexibility, improved reliability, and simplified problem diagnosis as compared to the in-stack version of the TN3270 Server. z/OS V1R8 is planned to be the last release of z/OS Communications Server which will support the in-stack version of the TN3270 Server. After z/OS V1R8 this capability will be removed from the product. In preparation for that change, customers should consider implementing the standalone TN3270 Server. For more information please refer to http://www.ibm.com/software/network/commserver/zos/

AnyNet is removed from z/OS Communications Server. AnyNet has not been enhanced in many years, and has been supplanted by Enterprise Extender, which has superior function and performance. The ability to define parallel EE Transmission Groups (TGs) using multiple SAPs is removed in z/OS V1.8. Parallel TGs defined in such a manner provide no benefit over single EE logical links.

There is a **new TCP/IP configuration parameter** that allows users to designate the source IP address to be used for outbound TCP connections based on the destination IP addresses or networks. In scenarios where outbound TCP connections from z/OS need to traverse multiple distinct networks, this enhancement can help make firewall administration easier, as users can permit traffic from a single, predictable IP address to traverse the firewall.

In z/OS V1.8, the **Firewall Technologies component of the Integrated Security Services element is removed.** Many Firewall Technologies functions have been stabilized for some time and can be replaced using comparable or better functions provided by Communications Server, notably IPSecurity. In addition, a functionally rich

downloadable tool replaces the IPSecurity and IP Filtering configuration GUI support. For more information, refer to z/OS V1R7.0 Migration.

Slide 25

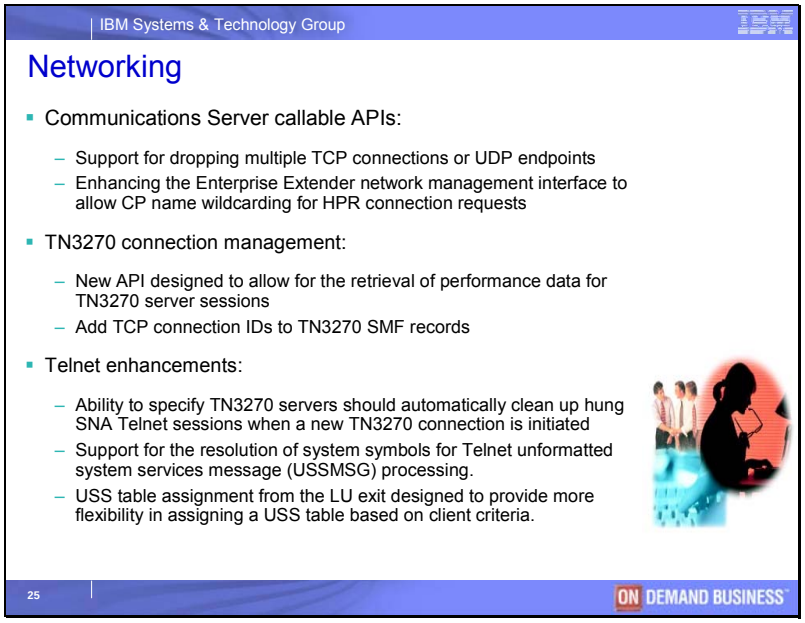

Slide 26

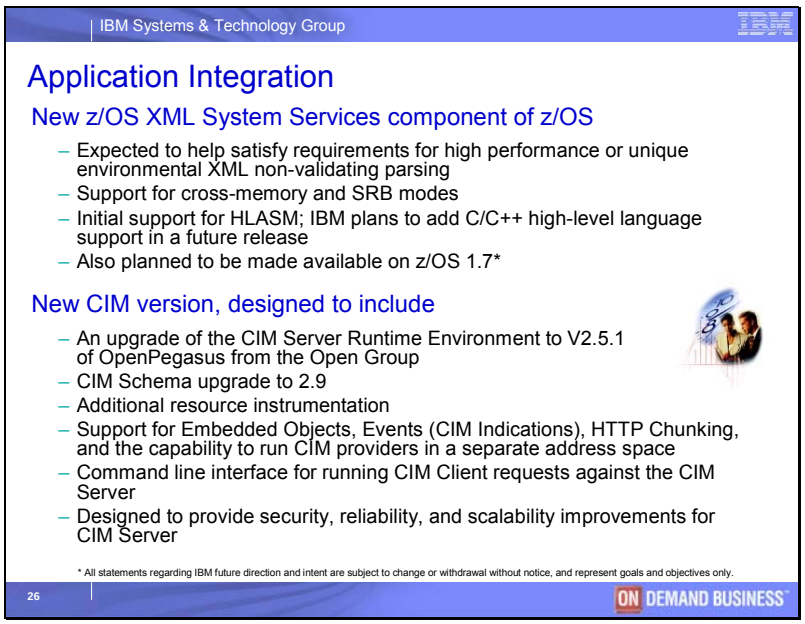

Application integration

In z/OS V1.8, there are several important functions that deliver further application integration support and support for industry and de facto standards:

- A new system component of z/OS, z/OS XML System Services
- Improvements in Unicode support
- BPXBATCH support for data sets
- CIM Server and schema upgrades
- The ability to trace transitions in and out of XPLINK in z/OS UNIX System Services (z/OS UNIX) programs
- Other improvements that are all intended to add new capabilities and help extend and tune your applications

New and enhanced z/OS V1.8 function includes:

- XL C/C++ has made the following performance and usability enhancements for z/OS V1.8:
  - XL C++ now provides support for the following set of C99 (ISO/IEC 9899:1999) features to aid in portability:
  - restrict qualifier
  - valid universal character name ranges
  - __func__ identifier for debugging assistance
  - XL C/C++ includes the following new compiler options and suboptions:
  - HGPR - to enable the compiler to use 64-bit instructions in 32-bit mode thereby potentially increasing the performance of 32-bit programs running on 64-bit hardware
  - HOT - to provide the ability to generate more highly optimized code

    - SPLITLIST - to split the IPA link phase listings into several files for easier viewing
    - LANGLVL(C99__FUNC__) - to assist with debugging code
    - The GONUMBER compiler option now supports 64-bit applications and generates line number tables for both 32-bit and 64-bit applications
    - The integrated CICS(R) translator is supported for use with CICS Translation Server V3.1 or later, to enable you to embed CICS statements in C/C++ source and pass them through the compiler without the need of an explicit preprocessing step permitting a more seamless operation of C/C++ within the CICS environment
    - Storage limitations are removed when optimizing complex applications using IPA, the IPA optimizer now runs in the 64-bit address space and no longer exhibits the storage limitations that restricted the use of the IPA optimizer during past releases
- The as command is available for HLASM assembler invocation to enable you to process assembler source files and invoke the HLASM assembler to produce object files

Note that z/OS V1.8 is planned to be the last release to include the C/C++ IBM Open Class„ (IOC) Dynamic Link Libraries (DLLs). For more information, refer to the z/OS statements of direction at http://www.ibm.com/servers/eserver/zseries/zos/zos_sods.html

- A **new LDAP server** for z/OS, IBM Tivoli Directory Server for z/OS, designed to improve performance, scalability, auditability, availability, and ease of use. This function is planned to be made available by 1H2007. The existing Integrated Security Services - LDAP Server will continue to be available in V1.8 in addition to the new IBM Tivoli Directory Server for z/OS. The IBM Tivoli Directory Server for z/OS delivers stronger affinity to z/OS platform features like Parallel Sysplex. Among the enhancements are:

- A new backend for small- and medium-sized directories designed to cache all directory entries in memory for better performance and use a z/OS UNIX System Services file as its backing store. This is made available in addition to the existing DB2-based backend. This enhancement is expected to help simplify setup and operation for small- and medium-sized directories.
- Automatic Restart Management (ARM) and TCP/IP restart support designed to help improve availability.
- Parallel Sysplex support designed to enhance synchronization of LDAP servers within a sysplex and allow a sysplex group to replicate with other LDAP servers.
- The creation of SMF records designed to improve LDAP auditability.

In V1.8, IBM introduces a new system component of z/OS, **z/OS XML System Services** (z/OS XML), which will be designed to deliver an optimized set of services for parsing XML documents. It is expected to be of use to IBM products, those from other software vendors, and customer middleware and applications having high performance or unique

environmental XML non-validating parsing requirements, such as the ability to run in cross-memory and service request block (SRB) modes.

Initial support delivers an assembler language interface. This new function is intended to satisfy a statement of direction made in Software Announcement 205-167 (RFA41468), dated July 26, 2005, and will also be available on z/OS V1.7. In a future release, IBM plans to add C/C++ high-level language support.

Unicode support is vital to enterprises with globalization imperatives; that is, to those for which storing data from different languages in a common format is important. z/OS V1.8 offers new collation tables needed by middleware applications. Their names are based on the Unicode Collation Algorithm (UCA), which conforms to the Unicode 4.0 standard. Using the new keywords can allow collation sequences for various locales to be used so that data stored in Unicode can be sorted using the appropriate collating sequence for each supported
language.

BPXBATCH program enhancements allowing STDOUT and STDERR DD statements to specify SYSOUT, partitioned, partitioned data set extended (PDSE), and sequential data sets in addition to z/OS UNIX files. Also, a new STDPARM DD statement is added that supports input from partitioned, partitioned data set extended (PDSE), sequential data sets, and z/OS UNIX files. The STDPARM data set supports parameter lists up to 64K bytes in length. When BPXBATCH is called from another program, the maximum size of the parameter list it will process without using STDPARM is increased to 32,754 bytes. Last, data sets specified using the STDENV DD statement may now specify a PDSE. This function is also available for z/OS V1.5 and higher with APAR OA11699.

z/OS V1.8 includes a new version of the **Common Information Model (CIM)**. This includes the upgrade of the CIM Server Runtime Environment to V2.5.1 of OpenPegasus from the Open Group, an upgrade of the CIM Schema to 2.9, and additional resource instrumentation. Key features of the new CIM Server for z/OS are the support for Embedded Objects, Events (CIM Indications), HTTP Chunking, and the capability to run CIM providers in a separate address space. In addition, a command line interface is provided to execute CIM Client requests against the CIM Server. Improvements are implemented in the areas of security, reliability, and scalability of the CIM Server.

RMF implements **eServer OS monitoring Stage 2 for z/OS.** Stage 2 includes:
A WBEM profile based on the IETF SLP protocol is intended to make self-discovery of the CIM-based monitoring service easier.
Support for CIM indications is designed to enable exploiters using standard CIM client applications to subscribe to asynchronous events.
Support for making additional RMF metrics available to CIM.

Language Environment is now designed to provide a new capability for **tracing transitions between executables compiled using the XPLINK compiler option and non-XPLINK executables.** This is expected to help application programmers more

easily diagnose the performance bottlenecks that might exist in mixed XPLINK/non-XPLINK applications.

The Program Management **Binder is now designed to provide Extended Relative Immediate Support**, and to help reduce the size of load modules and program objects with new options you can specify. The Program Management Binder is now designed to: Provide Extended Relative Immediate Support. The binder is now designed to further exploit the use the relative immediate instructions. These instructions can now be used with external references when creating object modules, and the binder is now designed to resolve those references between object modules.

Support new options for reducing the size of load modules and program options. The binder is designed to reduce the size of load modules and program objects by removing information unnecessary for execution when you specify the new STRIPCL and STRIPSEC options. STRIPCL specifies that removable classes be removed, while STRIPSEC specifies that unreferenced sections be removed.

Support for the **flockfile() family of functions** in Language Environment. The flockfile(), ftrylockfile(), funlockfile(), getc_unlocked(), getchar_unlocked(), putc_unlocked(), and putchar_unlocked() will be implemented in a way intended to comply with their description in the Single UNIX Specification Version 3 (SUSV3) standard.

There is support for the **/etc/inittab** file that is used on other UNIX systems to start and restart daemons. This is intended to allow you to identify system processes to be started during system initialization that should receive additional system management by z/OS UNIX System Services.

The **dbx debugger** is designed to support the complex and packed decimal data types used by applications that demand a higher degree of mathematical precision.

ISPF Edit informational messages are made available to ISPF Edit macros. As in the past, in ISPF Edit interactive mode, the message identifier, short message text, and long message text are displayed on the screen. This change makes the same information available to non-interactive macros in order to help allow them to handle more conditions.

Additional locale support for the Euro currency symbol is provided for countries that have recently joined the European Union: Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia, and Slovenia. As it does for other locales, this support allows the Euro currency symbol to be displayed and printed as the national currency symbol in these countries. This function is also available with APAR PQ99282 for z/OS V1.4 and higher.

RMF Spreadsheet macros are extended to analyze zAAP usage.

The Software Configuration and Library Manager (**SCLM**) component of ISPF includes the following enhancements:

- The ability to specify a language description when using the SPROF (SCLM Edit Profile) command and the SCLMINFO service to define an SCLM language
- Additional information about return codes in error messages generated by the COBOL parser (FLMLPCBL)

JES3 is enhanced to allow applications using the **SYSOUT Application Programming Interface (SAPI)** to retrieve SYSOUT data sets with READ access to the appropriate profile in the JESSPOOL class. Currently, UPDATE access is always required.

Slide 27

Application Integration

New LDAP server (1H2007)*

- Designed to improve performance, scalability, auditability, availability, and ease of use, and planned to include:
  - z/OS UNIX file-based backend for small- and medium-sized directories with memory caching for all directory entries, to help simplify setup and operation of small- and medium-sized directories
    - In addition to existing support for DB2-based backstore
  - ARM and TCP/IP restart support
  - Parallel Sysplex support designed to enhance synchronization of LDAP servers within a sysplex and allow a sysplex group to replicate with other LDAP servers
  - SMF records designed to improve LDAP auditability.

Session 1726 – z/OS LDAP Feature and Security Update

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

Slide 28

Application Integration

- Unicode collation tables:
  - New needed by middleware applications added
  - Based on the Unicode Collation Algorithm (UCA), which conforms to the Unicode 4.0 standard
- Improvements for Language Environment's Preinitialized Environments for Authorized Programs:
  - Support for changing the number of system-managed environments dynamically
  - Improved failure data collection capabilities and more recovery options
- inittab support in z/OS UNIX for Java recovery
  - Provide de facto standard recovery methods
- XPLINK transition tracing
  - Can help optimize application performance
- Additional locale Euro support for EU countries
  - Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia, Slovenia

Slide 29

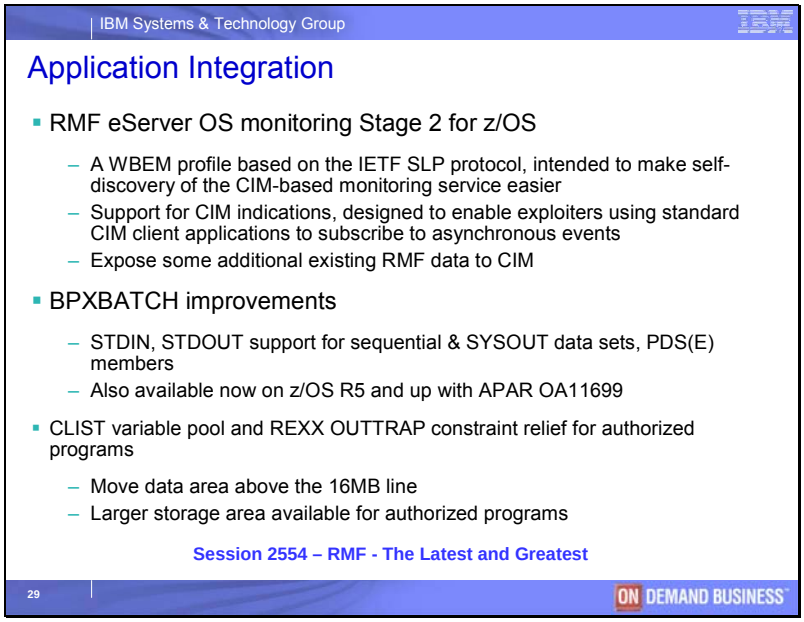

Slide 30

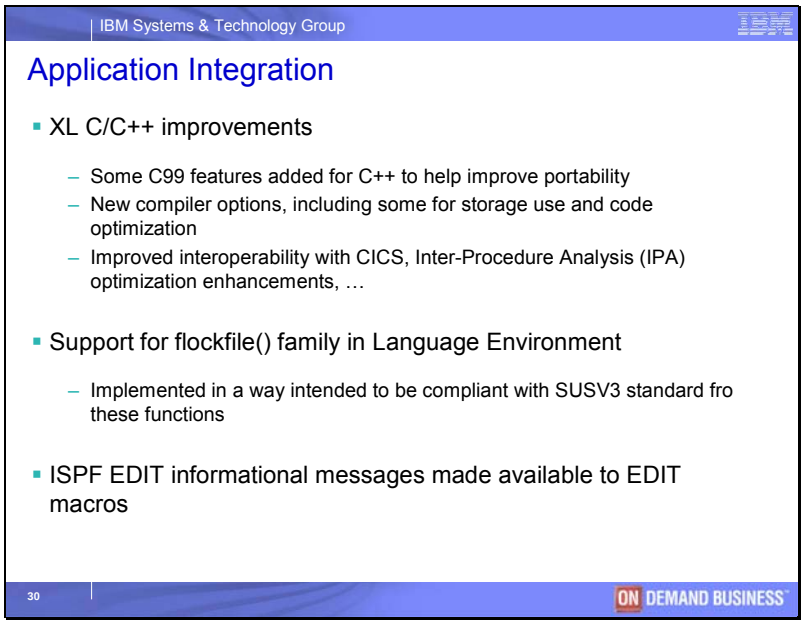

Slide 31

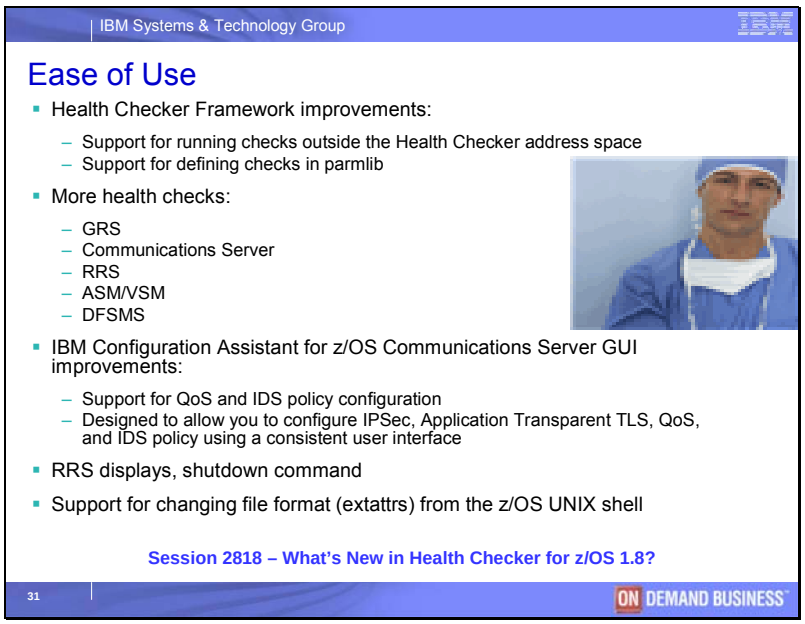

Ease of use

Ease of use is a continued focus area in z/OS V1.8. There are significant improvements delivered in the IBM Health Checker for z/OS framework and also additional checks, a number of Hardware Configuration Manager (HCM) and ISPF enhancements, and usability extensions to Infoprint Server's Web-based print management interface.

These and other usability improvements are:

The **IBM Configuration Assistant for z/OS Communications Server** GUI, introduced in z/OS V1.7 as the z/OS NSCA (Network Security Configuration Assistant), is extended to include support for QoS (Quality of Service) and IDS (Intrusion Detection Services) policy configuration. This support allows an administrator to configure IPSec, Application Transparent TLS, QoS, and IDS policy using a consistent user interface. In the past, QoS and IDS were configured with separate GUIs. These changes are expected to simplify the QoS/IDS configuration tasks by applying the Configuration Assistant for z/OS Communications Server concepts to QoS and IDS. Also, the GUI is designed to generate the policy agent configuration file for IDS policy, as described in Security.

In z/OS V1.8, there are significant improvements for **health checking**. The enhanced IBM Health Checker for z/OS framework makes it easier to write checks, and to provide improved parmlib, parsing, and display support for checks. In addition, there are a number of new checks for Communications Server, GRS, storage management (ASM), RACF, BAM, DFSMS, and Resource Recovery Services (RRS).

The following are **Hardware Configuration Manager (HCM) enhancements:**

- A new function is designed to export and import I/O definition files (IODFs). This improves processing time and helps make it easier to send IODFs to IBM for problem determination.
- A new function is designed to save the layout of a controller in a physical description file (PDF) so it can be used as a model for new controllers.
- Another new function allows you to locate objects that contain specific values in user fields, filter via wildcards, and select multiple objects for deletion in the Edit dialog.
- There is support for the Copy, Add like, and Repeat functions that are provided by Hardware Configuration Definition (HCD). This support enables you to define complex objects more quickly and eliminate the need to switch back to HCD to use the function.
- A new function is intended to allow you to compare HCM configuration files and get HCD IODF Compare Reports via an HCM dialog, and provide an additional means to check whether the changes that have been made were those intended.
- The automatic generation of entries in the activity log file of an IODF enhances the usability of the activity log file.
- The generation of cable labels includes PCHID information for channel paths. Switch port names are defaulted to the connection data for each port. These functions are intended to reduce the need for manual changes.
- The ability to hide connections in the HCM diagram helps allow you to tailor a cropped configuration view for documentation purposes.
- Hardware Configuration Manager (HCM) provides performance data integration with a Web interface of the RMF Distributed Data Server (DDS) for selected objects of the HCM diagram. The displayed RMF Monitor III online performance data enables you to detect performance bottlenecks and provide faster problem resolution.

The following are **ISPF improvements:**

- Support has been added for SuperC Compare and Search-For functions on the Data Set List Actions panel.
- Enhancements that help make it easier to search for members:
- A new member list primary command, FILTER, is designed so you can filter member lists using member attributes. The command is designed to allow you to use it repeatedly to refine the member list.
- A new option on the SRCHFOR command enables you to specify that member list filtering display only those members containing a search string.
- The ISPF workstation connect program WSCON is designed to auto-discover the IP address of the connected TN3270 workstation and use this address to establish a session with the Work Station Agent. This improves usability, because you will no longer need to be aware of your IP address or enter it on the Initiate Workstation Connection panel to establish a session.

- ISPF client/server code is converted to use the IBM C/C++ run-time libraries, which can help reduce the number of C run-time libraries in use on your system. Also, some performance improvement is expected.
- ISPF is designed to help ensure that, when a member having one or more aliases is renamed or deleted, aliases are updated to point to the new name or deleted. This helps prevent the inadvertent creation of "orphaned" alias entries.
- Data Set List line commands support for tape data sets is enabled. DFSMSrmm supports this by enabling fast path directly into the relevant part of the DFSMSrmm dialog.

Enhancements to **Infoprint Server's Infoprint Central** Web-based print management GUI are designed to let you:

- Display real-time status information for a TCP/IP-connected printer, including online/offline status, or how much paper is in the printer's input and output bins
- Turn a TCP/IP-connected printer online or offline, and reset the printer
- Stop an IP PrintWay job from printing without canceling the job
- Identify by user ID which Infoprint Central user performed a certain action on a print job, printer, or NetSpool LU
- Obtain details about network connectivity problems using the TRACEROUTE command for faster problem resolution

RMF enhancements designed to support ease of use include:

- The RMF postprocessor is designed to honor duration intervals for overview processing in the same way as for standard intervals. This support is intended to allow you to generate meaningful postprocessor trend reports over long periods and use the spreadsheet macros to plot charts for multiple days and even weeks.
- RMF Monitor III displays LRU status VSAM RLS activity reports, which enables you to distinguish between buffers below and above the 2-GB bar.

System Display and Search Facility **(SDSF) displays zAAP utilization** in addition to system CP utilization on the DA screen.

Recoverable Resource Services **(RRS) enhancements:**

- MVS DISPLAY commands include RRS status information. This makes it easier to implement automatic alerts and capture information in syslog.
- A new SHUTDOWN command is designed to allow you to end RRS instead of doing a CANCEL. This helps avoid unnecessary abnormal termination and provide clearer shutdown messages.

Slide 32

Ease of Use

- ISPF improvements
  - Display and process z/OS UNIX files and directories
  - Data Set List (OPT3.4) support for SuperC and Search-for
  - Better member list searching
  - Autodiscovery of IP address for ISPF workstation GUI
  - Improved alias processing for copy, move, delete
  - Tape data set support in DSLIST
- SDSF display of zAAP utilization
- RMF postprocessor & Monitor III improvements
- HCM enhancements
  - Literally, too many to list!
  - Some are: IODF import/export; saving layouts for modeling; Copy, Add Like, and Repeat support; and IODF compare

Slide 33

Enterprise-Wide Roles

- EWLM/WLM service class correlation
  - WLM designed to accept classification when same-named service class is passed from EWLM
- New DFSMSrmm CIM agent support
  - Designed to support the creation, change, and deletion of volumes and data sets
  - Complements the query and display functions in z/OS 1.7
- zAAP reporting for EWLM
- Infoprint Central GUI improvements
  - Real-time status information display
  - Change online/offline status, reset a printer
  - Stop a print job without canceling it
  - Auditability enhancements
  - TRACEROUTE support

Session 2712 – Infoprint Server V1R8
Session 3040 – DFSMSrmm for z/OS 1.8: A Preview
Session 3042 – WBEM/CIM Interface to DFSMSrmm

EWLM support is planned to be added to z/OS WLM, to accept the classification of work from EWLM, introduce a high-performance EWLM instrumentation option via WLM Enclave Services, map WLM execution delay monitoring services to appropriate Application Response Measurement (ARM) services, and provide a consistent management of process entitlements on both the z/OS-centric world and the

heterogeneous world controlled by EWLM by introducing new WLM system service classes.

WLM will be designed to include zAAP data in the CPU using and delay samples as well as in the CPU service times reported to EWLM for processes using zAAPs.

Enhancements to Infoprint Server's Infoprint Central Web-based print management GUI will be designed to:

- Display additional printer status information, such as paper levels in bins
- Turn printers online and offline and reset them directly from Infoprint Central
- Provide the ability to stop an IP PrintWay print job without canceling it
- Allow selection of log messages by hour and also by day
- Provide an audit trail for GUI actions into job and printer logs
- Force error messages to open when an action is performed on the Summary table
- Expand the ping action to include the TRACEROUTE command

Slide 34

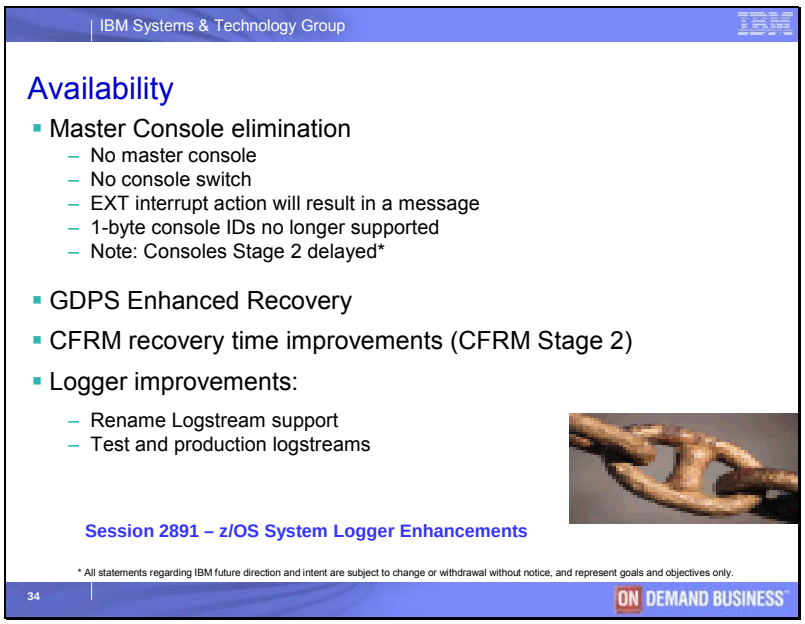

Availability

z/OS, in conjunction with IBM System z (TM) servers, continues to help address requirements for high availability. z/OS V1.8 is designed to extend DFSMS fast replication to the data set level on 2105 (Enterprise Storage Server) and later storage controllers; offers improved recoverability for the z/OS UNIX System Services (z/OS UNIX) Byte Range Lock Manager when a system failure occurs in a sysplex; improves the Coupling Facility Resource Manager (CFRM); provides further improvements to Consoles processing and System Logger; and can help make significant improvements in GDPS HyperSwap and CF duplexing recovery. These include:

In z/OS V1.8, **DFSMShsm fast replication** support extends the full-volume dump function to allow dumping of fast replication backup copies to tape. The dump copies can be created by command or by automatic dump. Fast replication backup versions can be recovered from DASD target volumes (as in previous releases) or from dump tapes.

In DB2 9 for z/OS, to further facilitate management of backup and recovery needs, the system-level utilities and the RECOVER utility will be enhanced to take advantage of new features provided by DFSMShsm Fast Replication support in z/OS V1.8.

The backup version being recovered can reside on disk or tape. When the backup version is recovered from disk, the recovery can be performed using fast replication or previously available copy methods. The combination of DB2 9 for z/OS Support with DFSMShsm Fast Replicate in z/OS 1.8 enables users to recover its backed up DB2 data extending back to any point in time subsequent to the earliest currently retained backup version (up to 50 backup versions allowed) and logs that have been stored.

In z/OS V1.8, the **z/OS UNIX System Services (z/OS UNIX) byte range lock manager (BRLM) is enhanced.** It is designed to allow applications that use byte-range locking to stay active when a system within the sysplex fails. This improves availability for those applications.

CFRM performance improvements. These enhancements are intended to enhance sysplex availability by significantly reducing I/O contention for CFRM couple data sets (CDS). This improves the performance of sysplex processes that use the CFRM couple data set. These processes include rebuilding Coupling Facility structures, Coupling Facility Duplexing establishment and failover, and connect/disconnect processing. When sysplex partitioning actions are taken, and when a Coupling Facility fails, Parallel Sysplex configurations are expected to recover significantly faster, improving sysplex availability.

In z/OS V1.8, the **single master console is eliminated**, and therefore it no longer can act as a single point of failure. The functions associated with the master console, including master command authority and the ability to receive messages delivered via the INTERNAL or INSTREAM message attribute, are able to be applied to any console, including EMCS consoles, in the configuration. In addition, the console switch function is removed, which removes a potential point of failure because you are able to define more than two consoles with master console authority that can be used should a primary console fail.

System Logger enhancements in z/OS V1.8 include:

- Support for renaming an existing log stream. This allows you to maintain the current data under a new name and enable new work to begin after defining a new instance of the log stream, and is expected to help reduce the impact of a damaged log stream by not requiring that the original data be deleted.
- Support for specifying different groups of log streams in the LOGR CDS. This function is intended to allow you to separate logger activity for test and production log streams. This can allow them to run in the same sysplex with less interference between them.

GDPS/PPRC enhancements are designed to help reduce the length and variability of failover recovery times by improving the consistency between secondary copies of data at the recovery site and duplexed structures in the Coupling Facility. New interfaces between the Cross-System Communications Facility (XCF) and GDPS enables "break duplexing" decisions for duplexed Coupling Facility structures to be made in a way that parallels the DASD "break PPRC mirroring" decision.

This enhancement preserves a usable copy of the duplexed Coupling Facility structures for recovery. It helps improve GDPS recovery time, enable consistent application restart times, and reduce the need for log-based recovery. The XCF support for this function is available for z/OS V1.5 and above with the PTFs for APAR OA11719.

In z/OS V1.8, there is support for a **new HyperSwap trigger** source type, I/O Timeout. In the past, HyperSwap function was invoked automatically when an error indication is returned for an I/O operation because of failures such as subsystem failures, boxed devices, or I/O errors. This support is extended to act on I/O timeouts detected by the Missing Interrupt Handler (MIH). The new function is designed to allow transactions to resume processing quickly on the secondary volumes. It can help deliver improved availability while helping reduce the need for operator intervention. HyperSwap is available with the separately-priced GDPS/PPRC and GDPS/PPRC HyperSwap Manager offerings.

z/OS V1.8 is designed to provide **improved z/OS UNIX System Services latch contention detection.** The output of the DISPLAY OMVS,Waiters operator command is enhanced to include a table of waiting threads and show file system latches held. These additional display capabilities help operators better identify the tasks that are waiting and the resources they are waiting for.

Slide 35

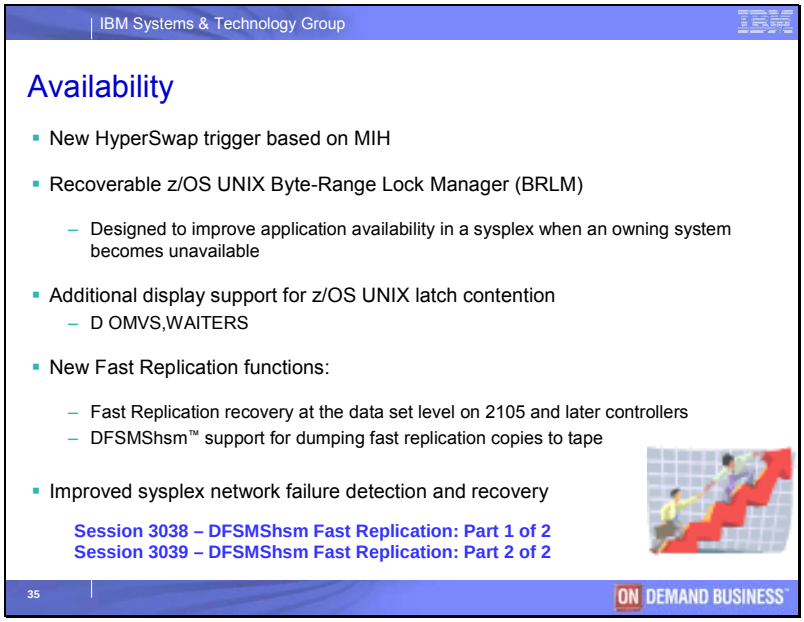

Slide 36

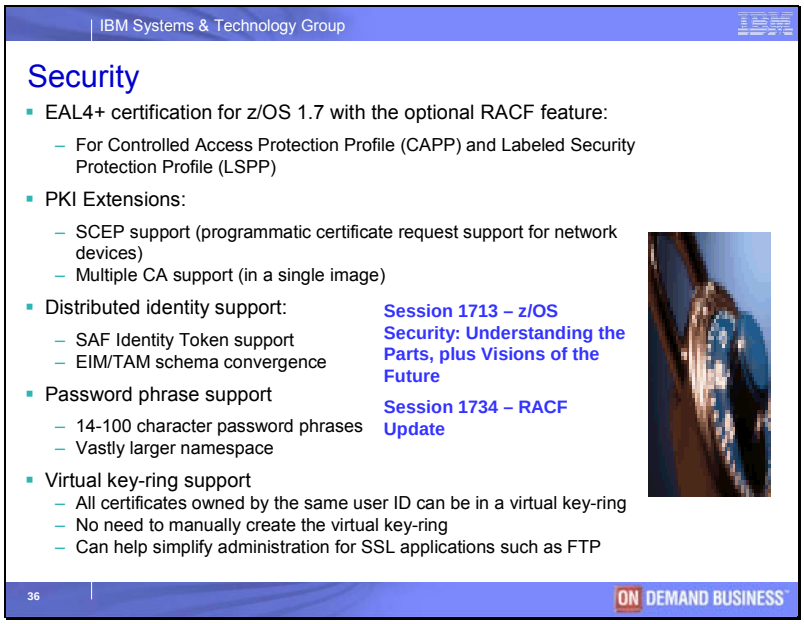

Security

In z/OS V1.8, z/OS continues to deliver industry leadership for security. Improvements that are all intended to help deliver the kind of security-rich environment that has made z/OS an industry leader include:

- Support for RACF password phrases from 14 to 100 characters in length
- Significant improvements to Identrus-certified support for digital certificates, including SCEP and multiple-CA support
- Improvements for tape data set protection using RACF or another external security manager
- Support for the Advanced Encryption Standard (AES) for IPSec

z/OS V1.8 is designed to include:

Improvements in the z/OS **System SSL** component are designed to help you secure communications for the industry-standard SSL and TLS protocols using the AES 128-bit hardware provided by the CP Assist Cryptographic Facility (CPACF) on System z9 servers. Additionally, z/OS System SSL now supports X.509 V3 certificates with SHA-256 and RSA encryption signatures in the z/OS System SSL key database. These certificates can now be used when using SSL and TLS.

Support for defining **Intrusion Detection Services (IDS) policies** in a policy agent configuration file as well as an LDAP server. This solution provides an IDS policy solution that is intended to be consistent with other policy types for those installations that do not have an LDAP infrastructure in place or that prefer using configuration files instead of LDAP.

RACF supports the use of passwords longer than eight characters, called password phrases. A password phrase is a character string that can comprise mixed-case letters, numbers, and certain special characters from 14 to 100 characters in length. Password phrases allow for an exponentially greater number of possible combinations of characters than do passwords. A user ID can have both a password and a password phrase. The same user ID can be used both for existing applications that accept an eight-character password and those that take advantage of the new password phrase infrastructure. This support is intended to help to improve system security and usability.

Public Key Infrastructure (PKI) Services improvements:

- Multiple certificate authority (CA) support are designed to allow more than one instance of the PKI Services daemon to run on a z/OS system. Also, new support is intended to allow you to establish multiple certificate authorities on a single image. Both new functions help improve the scale and availability of PKI Services management and fulfillment services on z/OS.
- PKI services is now designed to support the Simple Certificate Enrollment Protocol (SCEP). SCEP support, which is also offered by other industry certificate authority (CA) software, allows SCEP-enabled clients, such as Cisco routers, to request certificates by sending messages to a CA using the HTTP protocol. The addition of SCEP support is designed to allow PKI Services to respond to SCEP messages and support both manual and automatic enrollment.

DFSMS support introduces **new options for securing tape data sets using System Authorization Facility (SAF).** These are designed to allow you to define profiles to

protect data sets on tape using the DATASET class without the need to activate the TAPEDSN option or the TAPEVOL class.

DFSMS also provides options you can use to specify that all data sets on a tape volume should have common authorization and that users are authorized to overwrite existing files on a tape volume.

Support is provided for the **Advanced Encryption Standard (AES) algorithm for IP Security with a 128-bit key length.** This algorithm replaces DES as the standard encryption algorithm. Infoprint Server documentation has been enhanced to describe how this can be used to encrypt print output between z/OS and distributed printers.

Support for **SAF Identity Tokens** provides exploiters with increased user accountability and auditability of resources by providing end-to-end auditing that tracks identities used for initial authentication and those used on the current platform.

There is **RACF support for virtual key rings.** This support treats the collection of all the certificates owned by one user ID, including the SITE and CERTAUTH reserved user IDs, as an independent key ring. The use of the CERTAUTH virtual key ring is intended to help eliminate the need to manually create multiple real key rings for SSL-enabled z/OS client applications such as FTP.

Slide 37

Security

- Support for defining IDS policy in a file
  - In addition to via LDAP
- Improved tape data set security administration
  - Can use DATASET class without activating TAPEVOL or TAPEDSN
  - Can specify that all data sets on the same tape should have common authorization
- IPSEC support for 128-bit AES

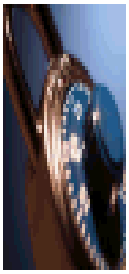

Session 3041 – Tape Security with DFSMSrmm

Slide 38

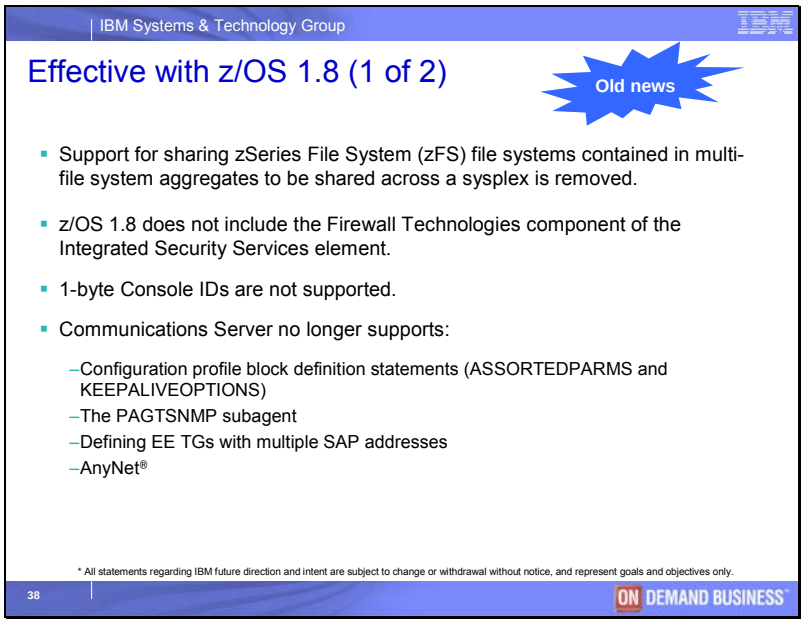

Statement of direction

IBM plans to take the following actions effective with the general availability of z/OS V1.8:

The APPC Application Suite is a set of common applications originally designed to enhance the value of SNA networks for end users. Since more full-featured alternative applications exist in modern integrated SNA/IP networks, z/OS V1.8 is planned to be the last release of z/OS Communications Server which will include the APPC Application Suite. After z/OS V1.8 the APPC Application Suite will no longer be shipped with the product, and will not be supported. However, note that APPC itself remains an integral part of z/OS Communications Server's SNA functions, and there are no plans to remove APPC from z/OS. For more information, refer to http://www.**ibm.com**/software/network/commserver/zos/

IBM plans to take the following actions in the future:

Preview — Group capacity limit on IBM System z9 EC and z9 BC : IBM plans to make it possible to define a logical partition (LPAR) group capacity limit on System z9 servers. This function will be designed to allow you to specify a capacity limit to be defined for each LPAR, and to define a group of LPARs on a server. This is expected to allow the system to manage the group in such a way that the sum of the LPAR capacity limits will not be exceeded.

When available, support of group capacity limit will be exclusive to System z9 EC and z9 BC.

Layer 3 Virtual MAC for z/OS and z/OS.e environments : When the associated OSA-Express function is available, the z/OS Layer 3 VMAC function will help to simplify the network infrastructure and facilitate load balancing when multiple operating system instances are sharing the same OSA port or Media Access Control (MAC) address. With Layer 3 VMAC, each operating system instance can now have its own unique "logical" or "virtual" MAC (VMAC) address. Instead of sharing the same physical MAC address, each TCP/IP stack and its associated IP addresses are accessible using their own VMAC addresses.

Layer 3 VMAC is an improved virtualization technique which is designed to dedicate a Layer 3 VMAC to a single TCP/IP stack which can help:

- Simplify network configuration
- Improve IP workload balancing
- Remove the dependency on Generic Routing Encapsulation (GRE) tunnels
- Allow WebSphere Application Server content-based routing to interoperate with z/OS in an IPv6 network
- Allow z/OS to use a "standard" interface ID for IPv6 addresses
- Remove the need for PRIROUTER/SECROUTER function in z/OS

Preview — When available, OSA Layer 3 VMAC is planned to be exclusive to z9 EC and z9 BC, and will be applicable to the OSA-Express2 features and to the OSA-Express features (Gigabit Ethernet LX and SX and 1000BASE-T Ethernet) when configured as CHPID type OSD (QDIO). OSA Layer 3 VMAC is required to exploit the support in z/OS V1.8.

Network Traffic Analyzer : When the associated Open Systems Adapter (OSA) function is available, the Network Traffic Analyzer function will help to provide a new diagnostic trace facility, allowing OSA packet trace records to be sent directly to the host operating system. Using Network Traffic Analyzer, along with the existing z/OS tools and service aids, z/OS system administrators will be able to monitor ("sniff") each OSA CHPID for the network traffic.

Network Traffic Analyzer is designed to extend the serviceability features of OSA, helping provide system programmers and network administrators the capability to control, capture, and format OSA hardware packet trace information.

Preview — When available, an enhancement to the QDIO architecture (OSA-Express Network Traffic Analyzer) will be designed to allow trace records to be sent to z/OS. This enhancement is planned to be exclusive to z9 EC and z9 BC, and will be applicable to the OSA-Express2 features when configured as CHPID type OSD (QDIO). OSA-Express Network Traffic Analyzer is required to exploit the support in z/OS V1.8.

QDIO Diagnostic Synchronization : When the associated OSA function is available, the QDIO Diagnostic Synchronization function will help provide system programmers and network administrators the ability to coordinate and simultaneously capture both software (z/OS) and hardware (OSA-Express2) traces. This function is designed to allow z/OS to

signal an OSA-Express2 feature (using a new Diagnostic Assist function) to stop traces and capture the current trace records.

Preview — When offered on z9 EC and z9 BC, QDIO Diagnostic Synchronization is planned to be applicable to the OSA-Express2 features when configured as CHPID type OSD (QDIO). The z9 EC and z9 BC QDIO Diagnostic Synchronization support is required to exploit the support in z/OS V1.8.

IBM plans to enhance the IBM Encryption Facility for z/OS (5655-P97) Encryption Services feature to use the OpenPGP standard, RFC 2440. This support will be designed to allow you to exchange an encrypted, compressed, and/or digitally signed file between your internal data centers using the Encryption Services feature in conjunction with your external partners and vendors who have an installed RFC 2440-compliant client running on z/OS and other operating systems. It is expected that IBM will implement the required functions of the OpenPGP standard, RFC 2440, that will allow Encryption Facility for z/OS to achieve compliance with the standard. The planned support includes, but is not limited to, symmetric encryption using AES and Triple-DES, asymmetric encryption of randomly generated symmetric keys using RSA and ElGamal algorithms, and working with DSA signatures.

On October 25, 2005, IBM announced the IBM Communication Controller for Linux for System z9 and zSeries V1.2, which provides an X.25 NPSI enablement interface. This function allows a software vendor to deliver support for an X.25 over TCPIP network. **It is IBM's intent to also release an IBM X.25 over TCPIP product** which uses this X.25 NPSI enablement interface. Together, the Communication Controller for Linux on System z and X.25 over TCPIP products can help you modernize your networking infrastructure for Communications Server on z/OS.

IBM intends to deliver enhancements to the consoles component in the future, not in 2007 as originally planned and announced in Software Announcement 206-039 , dated February 28, 2006. These enhancements will be intended to provide enhanced reliability and availability of the sysplex-wide management of the operator console configuration by reducing the amount of console configuration data passed around the sysplex under serialization. This change is expected to improve overall sysplex recovery time and reduce serialization delays during system startup and planned system shutdown activities. Increasing the maximum number of MCS and SMCS consoles that can be defined and active in a configuration from 99 per sysplex to 99 per system in the sysplex.

For more information, and for all previously announced statements of direction affecting z/OS V1.8 and future releases, visit
http://www.**ibm.com**/servers/eserver/zseries/zos/zos_sods.html

These statements represent current intentions of IBM.

Any reliance on these statements of direction is at the relying party's sole risk and will not create any liability or obligation for IBM. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

Slide 39

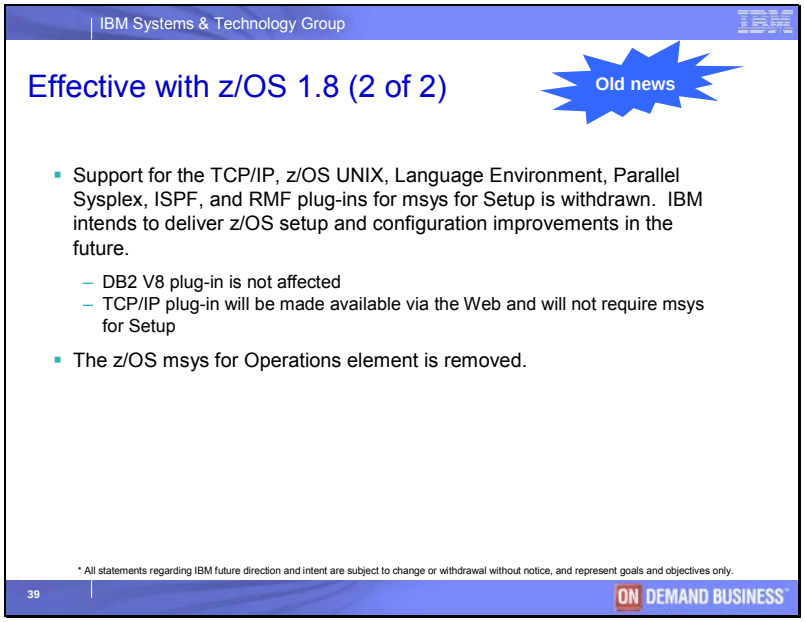

Slide 40

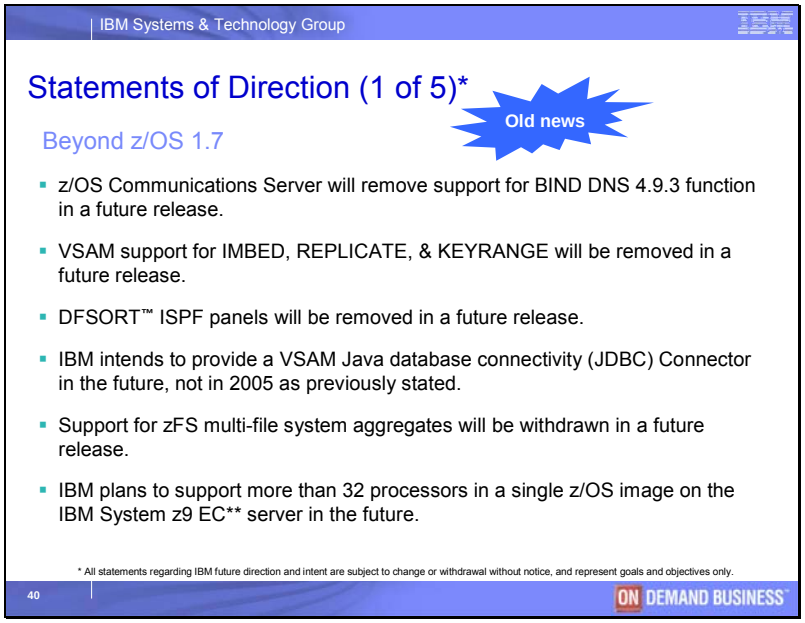

Slide 41

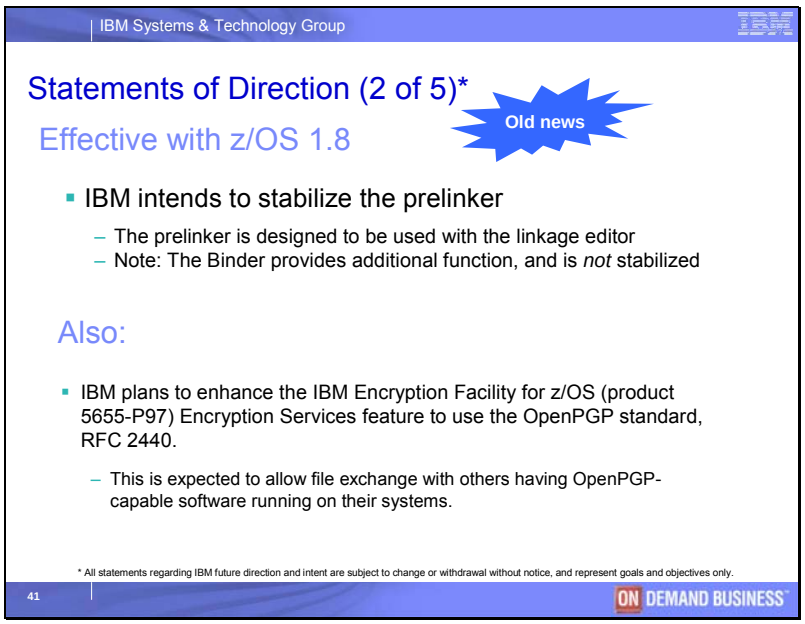

Statement of direction:

IBM plans to take the following actions effective with the general availability of z/OS V1.8:

IBM intends to stabilize the prelinker. The prelinker was designed to process long names and support constructed reentrancy in earlier versions of the C compiler on the MVS (TM) and OS/390 (R) operating systems. The prelinker, shipped with the z/OS C/C++ run-time library, provides output that is compatible with the linkage editor, shipped with the binder. The binder is designed to include the function of the prelinker, the linkage editor, the loader, and a number of APIs to manipulate the program object. Its functionality delivers a high level of compatibility with the prelinker and linkage editor, but provides additional functionality in some areas. For more information on the compatibility between the binder and the linker and prelinker, refer to *z/OS MVS Program Management: User's Guide and Reference*.
Further enhancements will not be made to the prelinker utility. Enhancements will be made only to the binder, the strategic tool for program object manipulation.

z/OS V1.8 is planned to be the last release that supports host communication between Hardware Configuration Manager (HCM) and Hardware Configuration Definition (HCD) elements via Advanced Program to Program Communication (APPC). Currently, for host communication either TCP/IP or APPC are supported. Starting with the z/OS release available in 2007, the host communication will be done exclusively via TCP/IP.

z/OS V1.8 is planned to be the last release to include the C/C++ IBM Open Class (IOC) Dynamic Link Libraries (DLLs). Application development support for the C/C++ IOC Library was withdrawn in z/OS V1.5. The run-time support (DLLs) for applications that

use the IOC Library is planned to be removed in the z/OS release available in 2007. Applications that are dependent on the IOC Library will not run starting with the z/OS release available in 2007. IBM has previously recommended that customers with application code that uses the IOC Library migrate to the Standard C++ Library. The publication *IBM Open Class Library Transition Guide* was published with z/OS V1.2 C/C++ as a reference for customers migrating their code from the IBM Open Class Library to the Standard C++ Library. You can get this guide by visiting **http://www-1.ibm.com/support/docview.wss?rs=32&org=SW&doc=7001423&loc=en-us**

Slide 42

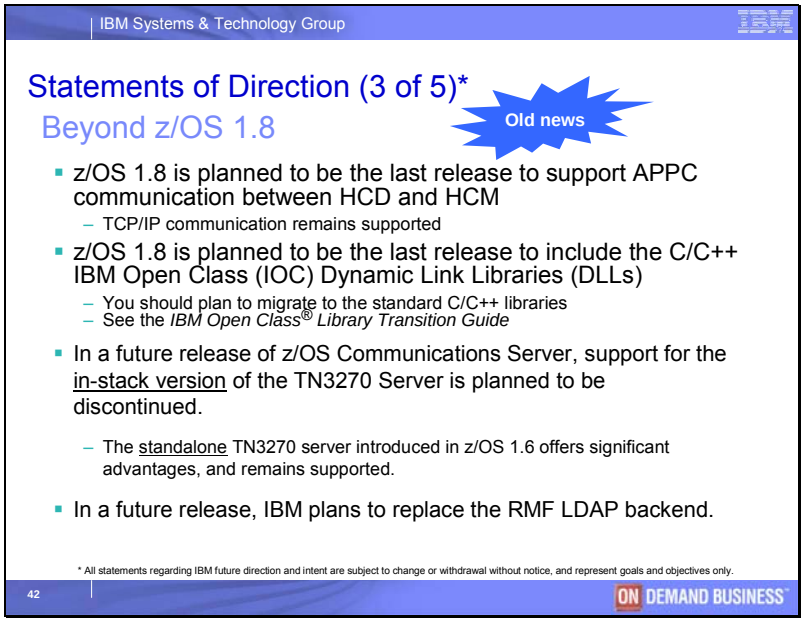

Slide 43

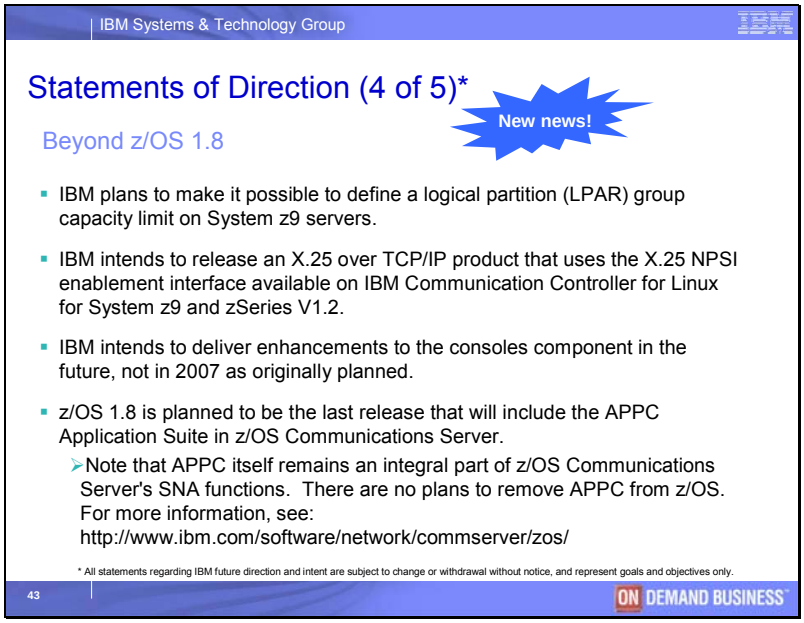

Slide 44

Statements of Direction (5 of 5)*

New news!

Beyond z/OS 1.8

- **Layer 3 Virtual MAC for z/OS and z/OS.e environments:** When the associated OSA-Express function is available, the z/OS Layer 3 VMAC function will help to simplify the network infrastructure and facilitate load balancing when multiple operating system instances are sharing the same OSA port or Media Access Control (MAC) address. With Layer 3 VMAC, each operating system instance can now have its own unique "logical" or "virtual" MAC (VMAC) address. Instead of sharing the same physical MAC address, each TCP/IP stack and its associated IP addresses are accessible using their own VMAC addresses.
- **Network Traffic Analyzer**: When the associated Open Systems Adapter (OSA) function is available, the Network Traffic Analyzer function will help to provide a new diagnostic trace facility, allowing OSA packet trace records to be sent directly to the host operating system.
- **QDIO Diagnostic Synchronization:** When the associated OSA function is available, the QDIO Diagnostic Synchronization function will help provide system programmers and network administrators the ability to coordinate and simultaneously capture both software (z/OS) and hardware (OSA-Express2) traces.

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

Slide 45

Related Sessions (1 of 2)

- Session 1633 – Improving the Health of Your z/OS Systems
- Session 1713 – z/OS Security: Understanding the Parts, plus Visions of the Future
- Session 1726 – z/OS LDAP Feature and Security Update
- Session 1734 – RACF Update
- Session 2540 -- WLM: The Latest and Greatest
- Session 2554 – RMF: The Latest and Greatest
- Session 2655 – JES2 z/OS 1.8 Product Overview & Review of Selected 1.7 Features
- Session 2712 – InfoPrint Server V1R8
- Session 2721 – z/OS 1.8 JES3 New Features
- Session 2818 – What's New in Health Checker for z/OS 1.8?
- Session 2870 – Migrating to z/OS, Part 1: Get Ready
- Session 2871 – Migrating to z/OS, Part 2: Get Set
- Session 2872 – Migrating to z/OS, Part 3: Go!
- Session 2884 – GRS: Forum at 6
- Session 2891 – z/OS System Logger Enhancements
- Session 2892 – One last drink from the fire hose: Migrating from z/OS R4 to R7
- Session 2896 - ShopzSeries and Serverpac/CBPDO - Simplified Ordering and Downloading of z/OS Products over the Internet!

Slide 46

Slide 47

Slide 48

System z Data Serving Roadmap

Today…

System z9 mainframes & z/OS combine for Industry Leading Performance, Scalability, Security, Availability and Virtual Partitioning capabilities with a focus on TCO improvements

DB2 Universal Database™ (UDB) for z/OS provides the premier high performance DBMS for high volume transaction workloads based on relational technology.

Security for z/OS and DB2 is designed to provide protection and advanced encryption capabilities with focus on recognized industry certifications

DB2 Content Manager for z/OS provides a foundation for managing, accessing, and integrating critical business information on demand.

WebSphere® Information Integration Platform for z/OS can help you readily access your pertinent information to support strategic business initiatives

DB2 Business Intelligence Solutions for z/OS provides cost-effective, rapid access to your core business data for warehousing and analytics

IMS™ & DB2 for z/OS Tools & Utilities provide integrated and intelligent management & monitoring of your IMS & DB2 applications and performance

IMS Version 9 extends and enhances your Industry-Leading IMS environment with increased connection & integration capabilities

Tomorrow…

DB2 V8 Exploitation of a new specialty processor (zIIP – System z9 Integrated Information Processor)
- types of eligible DB2 for z/OS V8 work, portions of which can be sent to the zIIP include:
- Network Connected Application processing using DRDA over a TCP/IP connection
- Data Warehouse Query processing that utilize star scheme parallel queries
- Select index maintenance in the DB2 Utilities (LOAD, REORG, REBUILD)

DB2 for z/OS Enhancements
- Unmatched infrastructure for supporting customer needs for regulatory compliance and auditing
- Continue to improve people resources efficiency
- Integrated XML
- DB2 Spatial Data Support
- OmniFind Search in DB2

Security for z/OS and DB2 Enhancements
- Common Criteria EAL Certifications
- Support of encryption capability (tape subsystem) with z/OS centralized key mgmt

z/OS Scalability & Resource Management Enhancements
- Enhancements in Device Addressability
- WLM Dynamic DB2 Bufferpool Management
- z/OS Global Mirror (XRC) and Global Mirror (asynchronous PPRC) consistency groups

WebSphere Information Integration Platform
- zSeries legacy data replication

DB2 and IMS Tools Focus Areas
- New and enhanced tools in recovery and performance
- Additional Support for SAP environments, Content Management & Business Integration

IMS Enterprise Support on z/OS
- XQuery Support
- Integrated Connect XML Adapter support

Future…

Further exploitation of zIIP

Additional DB2 for z/OS Enhancements
- Improve usability and reduce complexity
- Java Language Integrated Query
- DB2 Table Scan acceleration via DS8000 processing

Ongoing Security Enhancements
- Common Criteria EAL4 Certification for DB2
- Support of encryption capability (disk subsystem) with z/OS centralized key mgmt

System z9 server and z/OS Futures
- Continued FICON® Enhancements
- Continued focus on Storage capability & capacity
- Heterogeneous File Sharing Support
- NFS V4 Client
- Continuous data protect
- Dynamic storage provisioning

DB2 CM for z/OS Advancements
- XML API and Web Services

WebSphere Information Integration Platform
- Expanded integration and process automation across and between platform components

DB2 and IMS Tools Focus Areas
- New function in change management
- Increased function to help meet compliance requirements
- Continued advancements in key functional areas

Additional IMS Futures
- Enhancing IMS XML and Web Services Connectivity
- Integrating Operations across Subsystems/Platforms
- Simplifying Installation and Management

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

Slide 49

Continued investment in skills for the IBM Mainframe community

- **Enhancing mainframe skills today**
  - Training for customer "in-house" skills
  - Global road shows and workshops assisting in customer system designs
  - Significant investment in Technical Skills Vitality initiative and certifications
- **Advancing toward goal of 20,000 skilled mainframe professionals in marketplace by 2010**
  - Explosive growth in IBM Academic Initiative for System z; over 200 participating institutions
  - Educating university and college students worldwide
- **2005 IBM Student Mainframe Contest**
  - Over 700 students from 77 schools!
  - http://www.developer.ibm.com/university/students/contests/mainframecontest.html
  - Student comments:
    - I have submitted Parts 1 and 2 and am eager to continue. Lets do this :D -Chris Peterman
    - I don't know if I'll make it any further- schoolwork is eating up the time I should be investing in mainframe…but I'm enjoying it more than I can admit in public. -Beth Boose
    - I'm one of the "no experience required" guys the contest was advertised for and the whole thing is really interesting to me. I really had a blast with parts 1 and 2 and am looking forward to logging in to part 3. -Brian Capps
    - I enjoyed this contest very much, makes me appreciate my elders…kidding. Thanks for the great contest, -Walter Weiss
- **Draft of *z/OS Basics* available:**

  http://publibz.boulder.ibm.com/zoslib/pdf/zosbasic.pdf

Slide 50

http://www.ibm.com/servers/eserver/zseries/about/charter/community.html

Slide 51

Slide 52

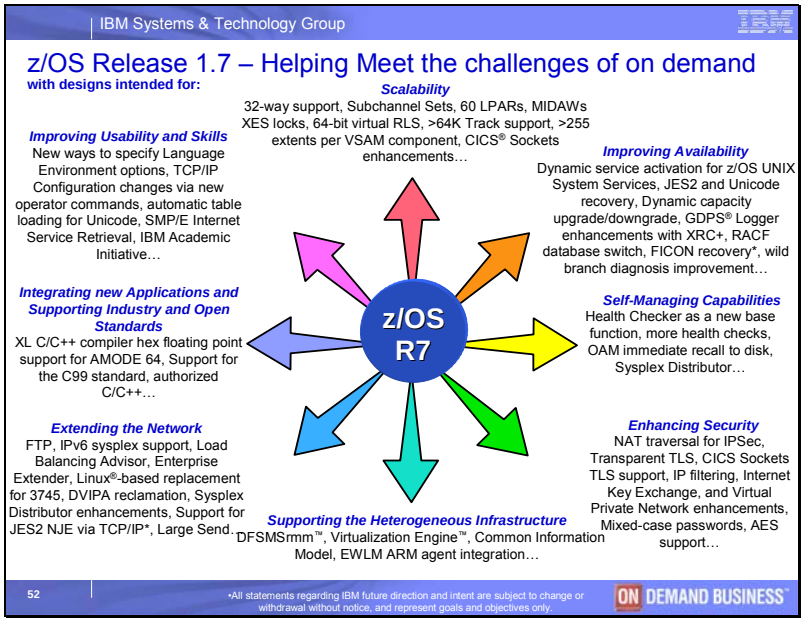

Slide 53

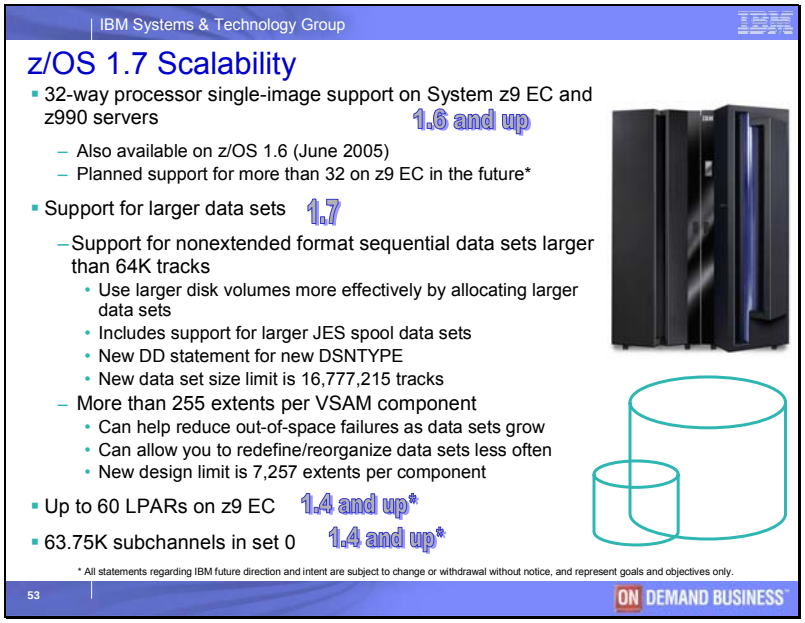

Scalability

z/OS continues to support robust vertical and horizontal growth. z/OS V1.7 leverages the new functions of the z9 EC and is designed to provide improved overall performance, increased I/O device addressability, increased I/O bandwidth, improved cryptographic performance when the Crypto Express2 feature is configured for Secure Sockets Layer (SSL) acceleration, and a new time synchronization feature, Server Time Protocol (STP), which is being previewed in this announcement. z/OS V1.7 also extends system limits in

many areas, including support for 32-way single-system images, larger sequential data sets, support for additional VSAM data set extents, more Cross System Extended Services (XES) locks per lock structure connector, 64-bit VSAM record-level sharing (RLS) support, and an increased number of DASD-only log streams. These improvements can help you support larger workloads for today's On Demand Business requirements.

Refer to Software Announcement 205-034 (RFA40707), dated February 15, 2005, for a description of these previously
announced functions:

- Up to 32 processors in a single image
- 64-bit VSAM record-level sharing (RLS)
- HCD support for larger I/O configurations
- A larger DFSMShsm tape table of contents (TTOC)
- DFSMSdss virtual storage constraint relief

IBM is once again doubling the number of Logical Partitions (LPARs). You now have the ability to define up to **60 LPARs** on the z9 EC, 15 LPARs per Logical Channel Subsystem, which may provide you even more flexibility to allocate hardware resources. With Processor Resource/Systems Manager (TM) (PR/SM) (TM) and Multiple Image Facility (MIF) you can share ESCON and FICON channels, ISC-3s, and OSA ports across LPARs. Support for up to 30 LPARs became available October 2003.

Support of up to 60 LPARs requires a z9 EC and is supported by z/OS, z/VM, z/VSE, TPF, z/TPF, and Linux. Refer to the *Hardware and Software requirements* sections of this [hardware] announcement. Introduction of 60 LPARs satisfies the Statement of General Direction in Hardware Announcement 103-142, (RFA38035) dated May 13, 2003.

Support for additional subchannels on System z9 EC servers: The z9 EC makes additional subchannels available. Previous servers reserved 1024 subchannels, making it possible to define a maximum of 64,512 devices. The z9 EC makes an additional 768 subchannels available, making it possible to define up to 65,280 devices for each z/OS LPAR.

63.75K subchannels: The z9 EC addresses a maximum of 64K-1 subchannels in subchannel set 0 (zero) with 1024 (1K) of these previously reserved for system use. IBM is making 768 of these 1K reserved subchannels available for customer use. The increased addressable storage this represents may be significant. For example, if you are using 3390 volume sizes and have 768 volumes of 54 GB/volume, this represents 41 terabytes (TB) of increased storage addressability (54 GB/volume * 768 volumes = 41 TB). In addition, the IBM TotalStorage DS8000 Series can be defined to attach 63.75K unit addresses. Now with 63.75K in the host, there is symmetry between the server and storage subsystem.

63.75K subchannels requires a z9 EC, is supported by all channel types, and is supported by z/OS and z/VM. Refer to the *Software requirements* section [of the hardware announcement] for more information.

In addition, z/OS V1.7 provides:

Support for large-format sequential data sets: Support is provided for nonextended-format sequential data sets larger than 64K tracks. This is in addition to previously provided support for extended-format data sets larger than 64K tracks. ISPF, SADMP, IPCS, AMASPZAP, DFSMSdss, and DFSORT support the new maximum data set size. JES2 and JES3 support using these larger data sets for spool. In addition, DFSMShsm and DFSMSrmm now allow using them for journal data sets, which can allow more processing between CDS backups. The new limit on the size of these data sets will be 16,777,215 tracks.

Both JES2 and JES3 are planned to support spool data sets up to the maximum supported volume size (65,520 cylinders). This can allow you to use full-volume spool data sets and exploit the new maximum volume size for spool volumes. A new operand of the DSNTYPE parameter of the JCL DD statement will be used to indicate that a data set larger than 64K tracks in size is to be allocated. This will avoid incompatibilities with programs that process Format 1 DSCBs, DEBs, or use previous levels of BSAM NOTE and POINT macros.

Support for a **greater number of extents for components of VSAM data sets.** The limit is raised from 255 extents per component to 7,257 extents per component. (This new limit is the product of the 59-volume limit on the number of volumes per VSAM data set and the 123-extent limit per volume.)

More than 1024 connections to DASDONLY log streams: The limit on the number of DASD-only log streams supported by System Logger on a single system has been raised from 1,024 to 16,384. This can allow you to move more work to an image without encountering Logger resource constraints.

Binder compression of program objects: The Program Management Binder can now compress program objects stored in PDSEs and z/OS UNIX System Services file systems such as zFS and HFS to reduce their size.

IFASMFDP buffers above the 16 MB line: The SMF dump program, IFASMFDP, is changed to use I/O buffers above the 16 MB line. This can allow the use of more output DD statements (for example, for splitting SMF records used for various clients) and an increased number of buffers. This can help improve the utility and performance of the SMF dump program.

VARY command processing improvements: VARY command processing is changed and is intended to reduce the duration of enqueues for the SYSIEFSD Q4 resource. This

can help improve the performance of sysplex-wide VARY command processing for large numbers of devices and can help prevent apparent hang conditions. Also, VARY OFFLINE processing is now done in parallel, which is designed to decrease elapsed time for commands used to change the status of many devices.

Protecting real storage below 16 MB and improving paging throughput: Storage management improvements are designed to help conserve real storage below the 16 MB line and to handle page faults more efficiently during periods of heavy paging. These changes can help prevent real storage shortages below 16 MB as workloads grow, and can help make the system more responsive when paging rates are high.

Standalone dump (SADMP) improvements: Standalone dump (SADMP) is improved to accommodate the larger amount of storage being dumped on today's larger systems. Better I/O buffering helps to make SADMP take less time, and the order in which data is dumped is changed to capture data in priority order.

Generalized Trace Facility (GTF) and Component Trace (CTRACE) external writer enhancements: The GTF external writer and CTRACE support writing to VSAM linear data sets, allowing for striping, wrapping support, and recording very large traces. This is expected to be especially helpful for Communications Server diagnosis, which relies heavily on GTF.

Increased summary dump buffer size: The summary dump buffer size is increased to accommodate the larger amounts of summary dump data required on today's systems. This is designed to help preserve enough data at the time the dump is initiated to include an accurate snapshot, while not requiring the system to be set nondispatchable while the dump is being taken; in turn, this can help improve first-failure data capture.

Note: When you are using the new IBM System z9 Application Assist Processor (zAAP), which provides a specialized z/OS Java execution environment, the total number of processors defined in a z/OS logical partition is the sum of general purpose processors (CPs), IBM System z9 Application Assist Processors (zAAPs), and IBM System z9 Integrated Information processors (zIIPs).

Slide 54

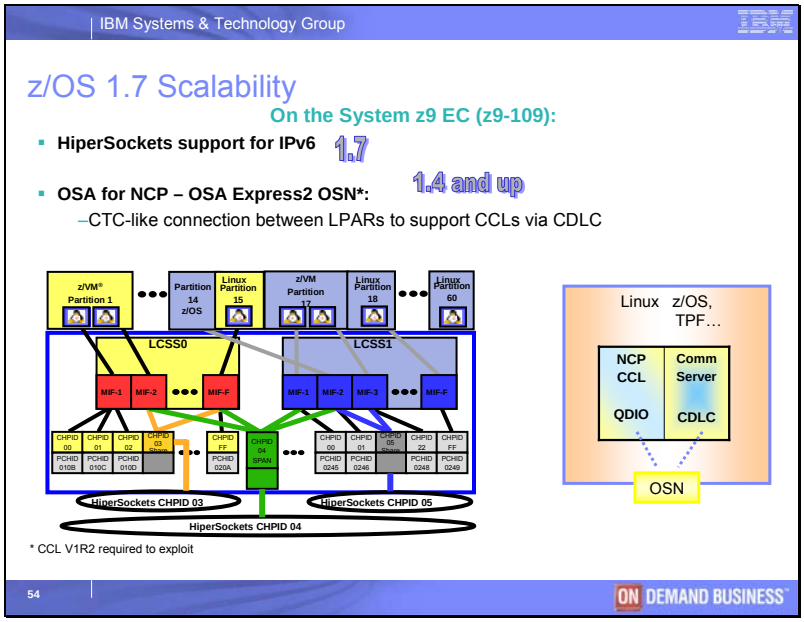

HiperSockets (TM) interfaces now support IPv6. This enhancement allows IPv6 communications between LPARs for z/OS TCP/IP stacks and Linux for zSeries TCP/IP stacks. It also expands IPv6 connectivity options between TCP/IP stacks in a sysplex when DYNAMICXCF is configured.

HiperSockets supports IP version 6: Internet Protocol Version 6 (IPv6) support is now being offered for HiperSockets (CHPID type IQD). IPv6 is the protocol designed by the Internet Engineering Task Force (IETF) to replace Internet Protocol Version 4 (IPv4) to help satisfy the demand for additional IP addresses.

IPv6 was introduced to expand the IP address space from 32 bits to 128 bits to enable a greater number of unique IP addresses in support of the proliferation of devices now connecting to the Internet, such as cell phones and PDAs.

The support of IPv6 on HiperSockets (CHPID type IQD) requires a z9 EC, and is supported by z/OS and z/VM. Refer to the *Software requirements* section [of the hardware announcement]. IPv6 support is currently available on the OSA-Express2 and OSA-Express features in the z/OS, z/VM, and Linux environments.

OSA CDLC support: OSA CDLC support is provided for z/OS and the IBM Communication Controller for Linux with APARs OA11238 and OA07875. This support is designed to allow z/OS to continue to communicate with an external network using channel data link control (CDLC) architecture, providing an alternative to an SNA network.

OSA-Express2 OSN - Open Systems Adapter for NCP: The OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features now have the capability to provide channel

connectivity from zNew operating systems to Communication Controller for Linux (CCL) with the introduction of the Open Systems Adapter for the Network Control Program supporting the Channel Data Link Control (CDLC) protocol.

OSA-Express2 OSN can help to eliminate the requirement to have any form of external medium (and all related hardware) for communications between the host operating system and the CCL image. Traffic between the two images (operating system and CCL) is no longer required to flow on an external Local Area Network (LAN) or ESCON channel; traffic can flow LPAR-to-LPAR. OSA-Express2 OSN provides an efficient method of communication, and is designed to create a secure and seamless integration of the host operating system and CCL.

If you continue to require SNA solutions that require NCP functions, you can now consider CCL as a migration strategy to replace your IBM Communications Controllers (374x). The CDLC connectivity option enables TPF environments to now exploit CCL.

The OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features provide connectivity with a new CHPID type OSN (Open Systems Adapter for NCP). Each host operating system that currently supports CDLC is expected to be able to utilize CDLC to CCL without changes to the operating system. OSA-Express2 OSN supports both SNA PU Type 5 and PU Type 2.1 channel connectivity.

Utilizing existing SNA support (multiple transmission groups), OSA-Express2 OSN support permits multiple connections between the same CCL image and the same host operating system (such as z/OS or TPF) image. The operating systems must reside within the same physical server (CPC) as the CCL image.

OSA-Express2 OSN:

- Is designed to appear to the operating systems as an ESCON channel connected to a 374x device type which exploits existing CDLC protocols
- Allows system administrators of the various operating systems to configure, manage, and operate their CCL NCPs as if they were running in an ESCON-attached 374x Communications Controller
- Enables NCP channel-related functions such as loading and dumping of the NCP
- Does not require external hardware (cables or switches)
- Allows multiple CCL images to communicate with multiple operating system images, supporting up to 180 connections (374x subchannels) per CHPID type OSN
- Can span Logical Channel Subsystems

The CCL image connects to OSA-Express2 OSN using extensions to the QDIO architecture.

OSA-Express2 OSN support requires a z9 EC and one of the OSA-Express2 Gigabit Ethernet SX, Gigabit Ethernet LX, and 1000BASE-T Ethernet features, and requires the port to be configured as CHPID type OSN, which can be configured on a port-by-port

basis. OSA-Express2 OSN is exclusively for internal communication, LPAR-to-LPAR. Concurrent LIC update is supported. OSA-Express2 OSN is supported by z/OS, z/VM, z/VSE, TPF, and Linux. Refer to the *Software requirements* section [of the hardware announcement]. Refer to the *Statement of General Direction* section [of the hardware announcement] for information on Communication Controller for Linux (CCL) support of OSA-Express2 OSN.

Slide 55

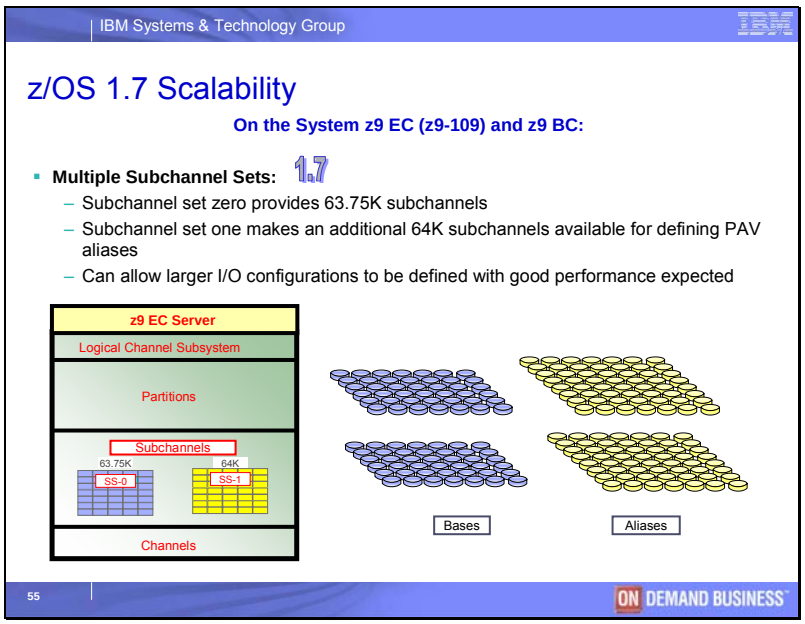

Multiple subchannel sets support: As systems and the corresponding need to increase the number of DASD volumes grow, there has been more pressure placed on the 64K device number, or unit control block (UCB) limit. The new multiple subchannel set support on z9 EC gives each single-system image the use of a second set of subchannels. z/OS V1.7 supports the use of a second subchannel set for defining Parallel Access Volumes (PAV) aliases on the z9 EC, which provides an additional 64K subchannels.

Moving PAV alias definitions to a second subchannel set can allow more devices to be defined. This new function provides relief from the 64K device limit by allowing PAV aliases to be defined without making device numbers unavailable for defining additional devices. EREP support is also provided to include the subchannel number for PAV alias devices.

Multiple Subchannel Sets: Multiple Subchannel Sets (MSS) is designed to provide an increased number of subchannels. Two subchannel sets are now available per LCSS and are designed to enable a total of 63.75K subchannels in set-0 and the addition of 64K-1 subchannels in set-1.

With the multiple subchannel set facility, one or two sets of subchannels may be configured to each LCSS where subchannel set 0 (zero) may have from 1 to 63.75K subchannels defined, and subchannel set 1 (one) may have from 1 to (64K-1) subchannels defined. z/OS will allow Parallel Access Volume Alias (PAV-alias) devices in the subchannel set 1 (one). MSS is designed to provide greater I/O device configuration capabilities for large enterprises.

MSS requires a z9 EC, and is supported by ESCON (CHPID type CNC), by all of the FICON features supported on z9 EC (CHPID type FC), by z/OS and Linux.

Slide 56

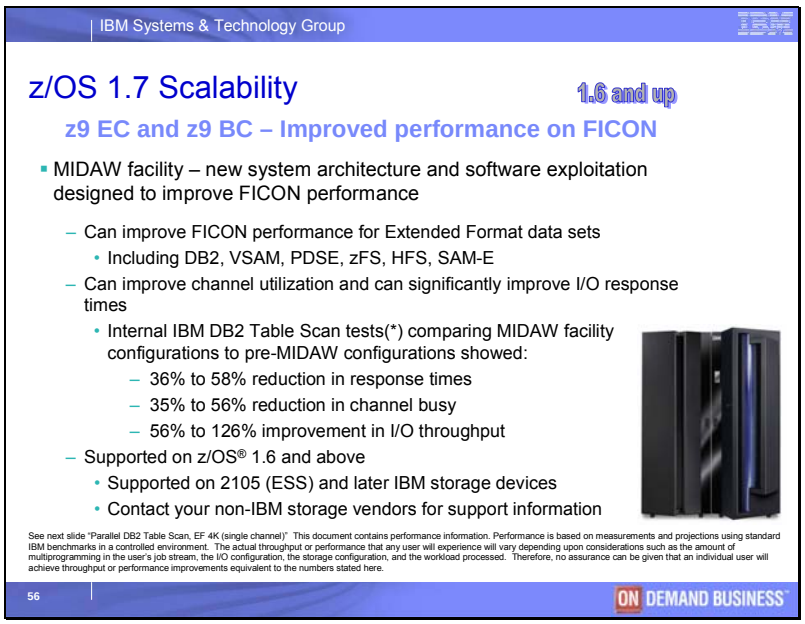

MIDAW (modified indirect addressing words) support: Indirect Addressing (IDA) provides virtual storage access for channel programs. The z9 EC implements a new function for channel programming, modified indirect addressing words (MIDAWs). MIDAWs can be used to move data over FICON and ESCON channels.

For FICON channels, this support can provide substantially better response time while increasing overall channel bandwidth. MIDAWs exploitation by z/OS is expected to improve performance for some DB2 table scan, DB2 sequential prefetch, BSAM, and extended-format data set operations by reducing system overhead for I/O requests on z9 EC servers, with no application changes.

Modified Indirect Data Address Word facility: z9 EC I/O architecture supports a new facility for indirect addressing, Modified Indirect Data Address Word (MIDAW) facility, for both ESCON and FICON channels. The use of the MIDAW facility, by applications that currently use data chaining, may result in improved FICON performance by reducing channel, director, and control unit overhead.
Applications that use the following may benefit: DB2, VSAM, Partitioned Data Set Extended (PDSE), Hierarchical File System (HFS), z/OS File System (zFS), SAM-E, and other datasets exploiting striping and compression. Applications using ESCON may also experience performance improvements.

The MIDAW facility requires a z9 EC, and is supported by ESCON (CHPID type CNC), and by the FICON features (CHPID types FCV and FC).

Slide 57

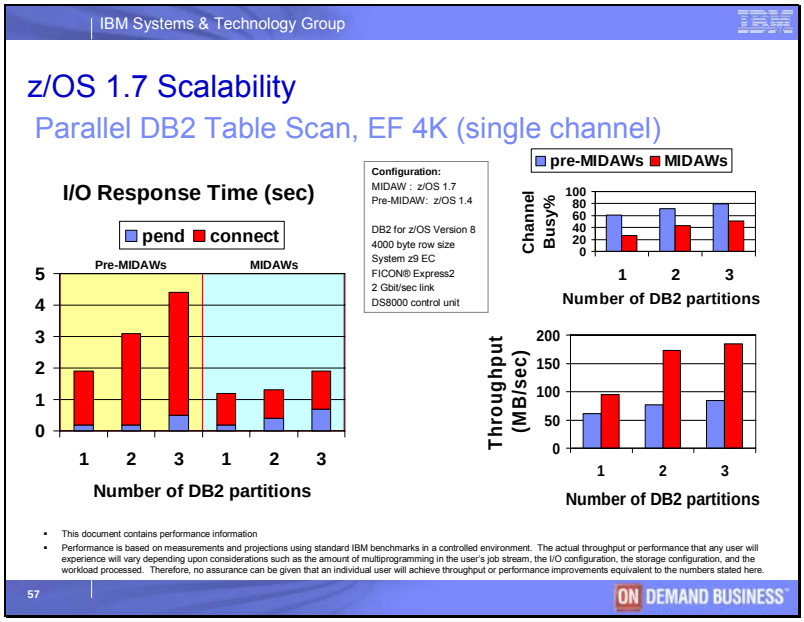

Slide 58

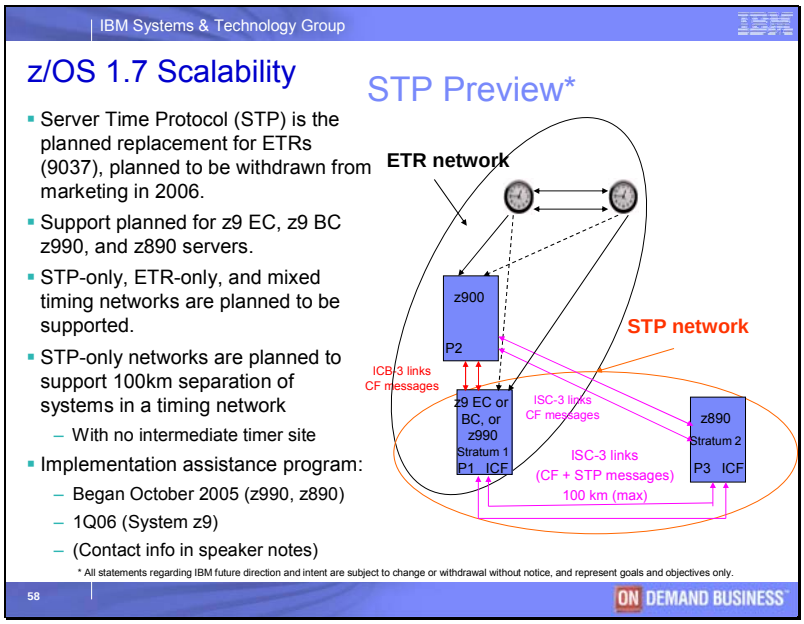

Server Time Protocol (STP) Overview:

Designed to provide capability for multiple System z9 and zSeries servers to maintain time synchronization with each other
Does not require the 9037 Sysplex Timer if all servers STP capable
Timing information transmitted over ISC-3 links (Peer mode), ICB-3 and ICB-4 links

Supports a multi-site timing network of up to 100 km (62 miles)
Allows a Parallel Sysplex to span up to 100 km
May reduce the cross-site connectivity required for a multi-site Parallel Sysplex
Can coexist with an External Time Reference (ETR) network (9037 based)
Mixed Timing Network
Designed to allow use of dial-out time services to set the time to international time standard (UTC) as well as adjust to UTC
Planned to be available as a feature on System z9 EC, z9 BC, z990 and z890
Prerequisites
z9 EC or z9 BC HMC
z/OS V1.7

Implementation Assistance Program (IAP):

Objective
- Accelerate the adoption of STP with IBM assistance

Assistance planned
- Consultation
- Review of migration plans
- Technical support

Availability dates:
- October 2005 (z990, z890)
- 1Q06 (z9)

Server Time Protocol (STP) preview: IBM plans to make available a new time synchronization feature, Server Time Protocol (STP), which is designed to provide the capability for multiple System z9 and zSeries servers to maintain time synchronization with each other. STP is planned to be the follow-on to the Sysplex Timer (9037-002). The Sysplex Timer and STP are designed to allow events occurring in different System z9, zSeries, and S/390 servers to be properly sequenced in time. Note that IBM intends to withdraw the Sysplex Timer Model 2 (9037-002) from marketing in 2006.

STP is designed for servers that have been configured to be in a Parallel Sysplex or a sysplex (without a coupling facility), as well as servers that are not in a sysplex, but that need to be time synchronized. STP is designed to allow timing information to be sent between servers and coupling facilities (CFs) over InterSystem Channel-3 (ISC-3) links configured in peer mode, Integrated Cluster Bus-3 (ICB-3) links, or Integrated Cluster Bus-4 (ICB-4) links.

STP is designed to:
- Allow clock synchronization for z9 EC, z9 BC, z990, and z890 servers and CFs without requiring the Sysplex Timer Support a multisite timing network of up to 100 km (62 miles) over fiber optic cabling, thus allowing a Parallel Sysplex to span these distances
- Potentially reduce the cross-site connectivity required for a multisite Parallel Sysplex

- Coexist with an ETR network
- Allow use of dial-out time services to set the time to an international time standard (such as Coordinated Universal Time (UTC)) as well as adjust to UTC on a periodic basis
- Allow setting of local time parameters, such as time zone and Daylight Savings Time
- Allow automatic updates of Daylight Savings Time

STP is planned to be available as a feature on z9 EC z9, BC, z990, and z890, and be supported by z/OS V1.7 (PTFs will be required to enable STP support). Implementation Assistance Program (IAP): IBM plans to make available an Implementation Assistance Program (IAP) to allow you to accelerate the adoption of STP with IBM's assistance. The assistance being planned for this program includes consultation, the review of your migration plans, and technical support. The IAP is planned to begin in October 2005 and be made available through 2006.

If you are interested in participating in the IAP, contact either:

Angelo Corridori – angeloc@us.ibm.com
Noshir Dhondy – noshir@us.ibm.com

Slide 59

Application integration
IBM continues to embrace open and industry standards to support your requirements for application portability. z/OS V1.7 provides C support for the ISO C99 standard, an ld utility in the Program Management Binder, additional open APIs and commonly-used device files in z/OS UNIX System Services, support for IEEE C fork() processing, and several dbx debugging enhancements. In addition, C/C++ support for authorized programs is provided, along with Program Management Binder and ISPF enhancements, and more flexible ways to specify Language
Environment options.

Hexadecimal floating point support for AMODE 64 C/C++ applications: z/OS XL C/C++ applications compiled using LP64 and FLOAT(HEX) are planned to be able to use the C/C++ run-time library math, numeric conversion, and formatted I/O functions that work with float, double, and long double data types. The initial C/C++ run-time library for AMODE 64 C/C++ applications on z/OS 1.6 provided floating point support for applications compiled using FLOAT(IEEE). This new support will complete IBM's planned floating point support within the C/C++ run-time library for AMODE 64 C/C++ applications.

Common Information Model (CIM) and Web-Based Enterprise Management (WBEM): z/OS V1.7 supports the Common Information Model (CIM), part of the Web-Based Enterprise Management (WBEM) initiative defined by the Distributed Management Task Force (DMTF). The CIM data model describes and accesses systems management data in heterogeneous environments. It can allow applications to measure system resources in a network with different operating systems and hardware. To enable z/OS to participate in cross-platform management from a common point of control, a subset of z/OS resources and metrics is mapped into the CIM standard data model to

allow its use for system management functions. In the release following z/OS V1.7, IBM plans to make the CIM services suitable for applications
that require high request rate processing.

XL C/C++: As of z/OS V1.7, the z/OS C/C++ compiler has been renamed to z/OS XL C/C++. The XL C and XL C++ compilers that are part of the C/C++ without Debug Tool optional priced feature of z/OS allow you to write code that follows the current ISO/IEC International Standards.

In z/OS V1.7, the XL C compiler is designed to support the latest ISO C 1999 (International Standard ISO/IEC 9899:1999) standard, also known as C99. The C99 standard offers a number of additional language functions to promote portability of C programs. The c99 command is now available, through the xlc utility, to invoke the compiler.

The z/OS XL C++ compiler is designed to support the International Standard for the Programming languages - C++ (ISO/IEC 14882:2003) and Programming languages - C++ (ISO/IEC 14882:1998). The LANGLVL compiler option and pragma are enhanced to support the available language levels that you can use to compile your C and C++ code. Suboptions are designed to provide conformance to the C99 standard, extended conformance to C99, conformance to the previous C89 standard, and extended conformance to C89. XL C/C++ offers you the flexibility to compile your code using the language level that meets your needs and is designed to improve usability and portability of programs across different platforms.
Note: The OS/390 V2.10 level of the C/C++ compiler is removed in z/OS V1.7.

ld utility: The Binder provides an ld utility with syntax similar to AIX's ld, with defaults that are general (not specific to Language Environment or C/C++) and with a mechanism for changing the defaults.

64-bit relative-immediate instructions (BRAS, BRASL, LARL): The Binder adds support for binding modules that include 64-bit relative-immediate instructions (BRAS, BRASL, LARL), which use external 3 symbols in their operands. The assembler added support for the BRAS, BRASL, and LARL instructions in z/OS V1.6. With z/OS V1.6, you could write and assemble, but not bind or run, programs using these
3 instructions. With z/OS V1.7, you can also bind and run such programs.

ISPF enhancements:

- Tracing support is added for file tailoring skeletons and for panels.
- ISPF now allows you to customize the default setting of the STACK option for the LIBDEF service. It sets a return code of 4 when the STKADD option is specified on a LIBDEF request and a LIBDEF has not been previously stacked. An advanced ISPF table editor can help you develop and debug applications that use ISPF tables.

Preinitialized environments for authorized programs: Preinitialized environments for authorized programs provides a Language Environment facility that supports the execution of Language Environment-enabled C/C++ code in an authorized environment.

More flexible ways to specify Language Environment options: Language Environment supports specification of run-time options in a file specified via DD statement. This makes it possible to specify Language Environment run-time options while avoiding the 100-character limit for the JCL PARM field. This change also provides a way to specify run-time options for IMS transactions when Library Routine Retention (LRR) is not used.

A new parmlib member, CEEPRMxx, can be used to specify Language Environment run-time options for the system. Operator commands are also provided that allow you to query and update the active run-time options for the system. This can simplify the management of Language Environment options, particularly in multisystem environments, and makes it possible to move Language Environment customization out of assembler language modules maintained using SMP/E usermods.

Specifying Language Environment options using CEEDOPT, CEECOPT and CELQDOPT modules will continue to be supported.

Additional open APIs for z/OS UNIX System Services: Additional open APIs for z/OS UNIX System Services (z/OS UNIX) provide native interfaces that can help make it easier to port IBM's and other vendors‘ applications on z/OS. These interfaces are added:

- pthread_getconcurrency()
- pthread_setconcurrency()
- pthread_setcancelstate()
- pthread_setcanceltype()
- pthread_testcancel()
- pthread_key_delete()
- pthread_sigmask()
- sched_yield()
- strerror_r()
- unsetenv()

The fork( ) C function is enhanced to work in a multi-threaded environment: According to IEEE Std 1003.1-2001, if a multi-threaded process calls fork(), the new process shall contain a replica of the calling thread and its entire address space, possibly including the states of mutexes and other resources. The intended use of fork() in a multi-threaded process is to run a new program, that is, invoke one of the exec() family of functions.

dbx scanner improvements:

- DBX - C++ typecast: The dbx scanner now recognizes C++ typecast expressions that contain the C++ "::“ operator, so that it is able to parse all valid C++ typecasts.

- Multiple view context: The dbx debugger supports contexts that contain multiple views that can be generated by the compiler/CDA.
- DBX - plugin architecture: In support of Java structural requirements, dbx has implemented the AIX dbx_plugin.h architecture on z/OS dbx. This support includes:
- The ability to read the contents of memory
- Notification when the user has issued a plugin command
- A "pluginload" command that allows users to load plugin libraries
- A "plugin" command that allows users to give commands to a loaded plugin

Addition of /dev/random and /dev/zero: The /dev/zero and /dev/random devices are added to z/OS UNIX System Services. /dev/zero is a character special device file. Data written to this file is discarded, and data read from this file is in the form of binary zeros. /dev/random is a device file that produces random bytes.

IBM XML Toolkit for z/OS, V1.8 (5655-J51) continues to provide enhanced C++ XML Parser and C++ XSLT Processor support, as announced in Software Announcement 205-114 (RFA41266), dated May 17, 2005. The toolkit provides the XML4C V5.5 XML parser and the XSLT4C V1.9 XSLT processor, based on corresponding Apache Software Foundation Xerces and Xalan technologies. IBM XML Toolkit for z/OS, V1.8 now includes versions of its libraries built using XPLINK, the extra performance linkage option. For more information, visit
http://www.ibm.com/servers/eserver/zseries/software/xml

Slide 60

z/OS 1.7 Integrating New Applications

- Support for authorized C/C++ programs in Language Environment
  - API to support APF-authorized C/C++ code and conforming assembler code
  - New authorized pre-initialization facility, CELAAUTH
  - Create environments for authorized code
  - Both system- and user-managed environments supported
  - Support for supervisor state, task and SRB modes, cross-memory mode
- Support for 64-bit hex floating point functions in C/C++ applications
- Binder support for 64-bit relative-immediate instructions
  - BRAS, BRASL, LARL
- ld utility in the Binder
  - Similar to AIX® ld
- C fork() enhanced to work in a multithreaded environment

Slide 61

z/OS 1.7 Integrating New Applications

Supporting Open & Industry Standards

- XL C/C++ ISO C99 Standard compliance for C programs
  - Complex and floating point APIs added
- More open APIs for z/OS UNIX System Services
- Added the /dev/random and /dev/zero devices to z/OS UNIX
- CIM support
- dbx scanner improvements
- XML Tool Kit for z/OS, V1.8
  - This is Release 8 of the tool kit for z/OS, not the z/OS 1.8 level of the tool kit!
  - Runs on z/OS Release 4 and up
  - C++ XML Parser and XSLT Processor support
  - Now available with libraries built using XPLINK
  - Note that in Release 7, Java XML Parser and Java XSLT Parser functions moved to:
    - Java SDK V1.3.1 level with PTF UQ77055
    - IBM SDK for z/OS Java 2 Technology Edition V1.4 (5655-I56) with PTF UQ77468
    - IBM 31-bit SDK for z/OS, Java 2 Technology Edition, V5 (5655-N98)
    - IBM 64-bit SDK for z/OS, Java 2 Technology Edition, V5 (5655-N99)

Slide 62

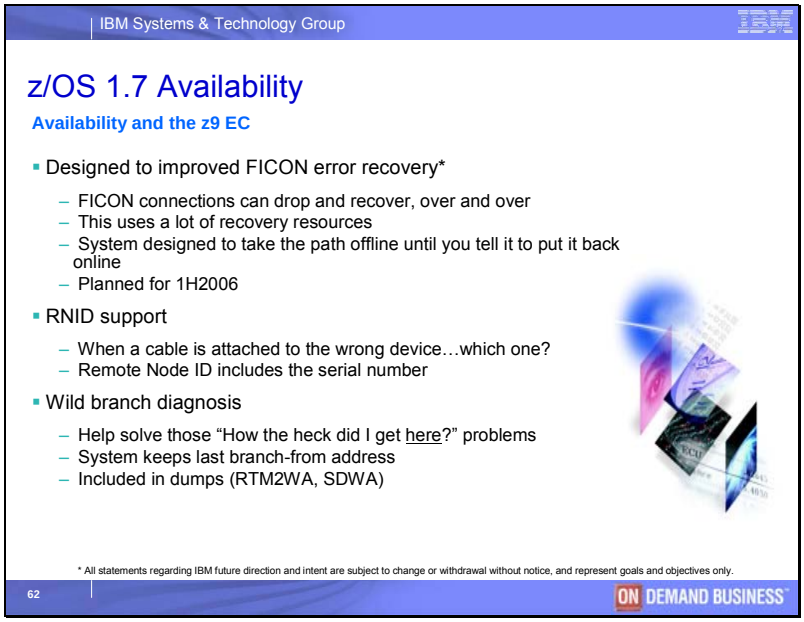

Slide 63

IBM Health Checker for z/OS is a new z/OS base function and provides a foundation to help simplify and automate the identification of potential configuration problems before they impact system availability. It compares active values and settings to those suggested by IBM or defined by your installation.

IBM Health Checker for z/OS consists of:

- The **framework**, which manages functions such as check registration, messaging, scheduling, command processing, logging, and reporting. The framework is provided as an intended programming interface to support writing new checks.
- **Checks**, which evaluate settings and definitions specific to products, elements, or components. Checks are provided separately and are independent of the framework. The architecture of the framework supports checks written by IBM, independent software vendors (ISVs), and users. You can manage checks and define overrides to defaults using the MODIFY command or the HZSPRMxx parmlib member.

Components and elements that have provided checks for z/OS V1.7 include consoles, cross-system coupling facility (XCF), Global Resource Serialization (GRS), real storage manager (RSM), resource recovery services (RRS), SDUMP, Security Server RACF, virtual storage manager (VSM), and z/OS UNIX System Services.

IBM-supplied checks may be integrated with the product, element, or component, or they may be provided as PTFs. Delivering checks in PTFs makes it possible to provide checks between releases. To easily identify checks that are provided as PTFs, you can use the Enhanced Preventive Service Planning Tool, available at:
http://techsupport.services.ibm.com/390/psp_main.html

You can identify checks by selecting a type of "Function" and a category of "Health Checker". Many of the checks are also supported on z/OS releases V1.4, V1.5, and V1.6; however, you should review the check PTFs for specific releases that are supported. SDSF for z/OS V1.7 provides support to make management of your checks easier with the new CK panel for the IBM Health Checker for z/OS. You can use the CK panel to display checks, attributes, and status, taking advantage of standard SDSF sort, filter, and arrange support. In addition, you can:

- Alter check attributes (such as status, interval, severity, category, and WTO descriptor)
- Browse check output for the most recent check
- Print check output or send it to a data set

The SDSF support will also be made available in PTFs for z/OS V1.4, V1.5, and V1.6.

IBM Health Checker for z/OS is also planned to be made available for z/OS and z/OS.e releases V1.4, V1.5, and V1.6 as a z/OS Web download from the z/OS Downloads page:
http://www.ibm.com/servers/eserver/zseries/zos/downloads/

OAM immediate recall to disk will implement the immediate recall of objects currently residing on removable media to DB2 for an installation-specified number of days. This can allow subsequent requests to read the objects from DB2 disk, rather than from a tape or optical volume. It can provide a significant performance improvement.

WLM and TCP/IP Sysplex Distributor will be designed to work together to route work to systems that are achieving response time goals. DB2 Connect will also use this new

function. In addition, routing will be designed to be sensitive to resource shortages (such as reaching MAXDBAT in a DRDA region).

Slide 64

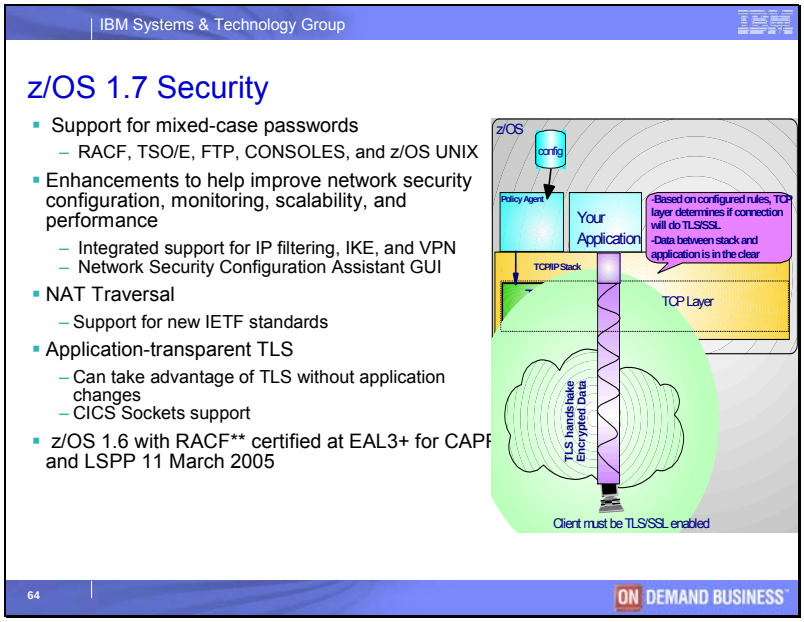

Mixed-case password support is planned to support lower- and mixed-case characters in passwords, in addition to uppercase characters. This increases the number of possible passwords considerably and can help improve interoperability. Support is provided in RACF, Consoles, FTP, z/OS UNIX, and TSO/E for mixed-case passwords. A RACF SETROPTS option allows you to enable or disable this function.

IPSec & NAT: Problems with network availability can occur when IPSec is used to protect traffic that traverses a system using Network Address Translation (NAT). New IETF standards address the interaction of IPSEC and NAT. This support is known as **NAT Traversal**. Communications Server improves availability by providing NAT Traversal support for a defined group of configurations.

A new, **application-transparent Transport Layer Security (TLS)** function in the TCPIP stack will provide TLS for TCP applications that require secure connections. Today, Transport Layer Security (TLS) is a network security protocol that must be managed at the application level. Applications requiring secure data connections must be modified and uniquely configured to manage the secure connection, which can result in high development and maintenance costs. Application-transparent TLS provided by the z/OS TCPIP stacks transparently provides TLS protected secure connections without requiring application change or awareness of the secure connection. This can help result in reduced application development and maintenance costs while improving the security of applications on z/OS.

Slide 65

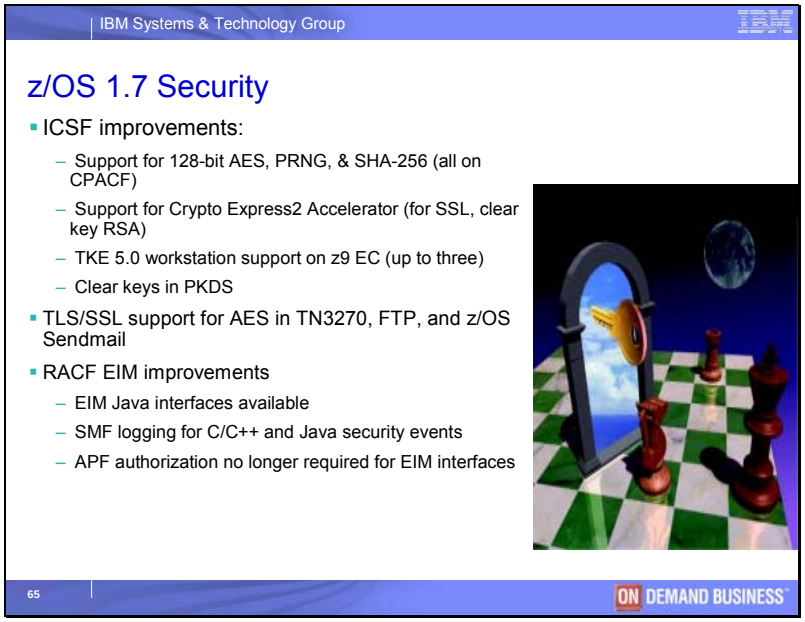

Security
Continued enhancements to security extend z/OS's leadership. In z/OS V1.7, significant improvements are made to RACF password processing, network security, auditability, public key infrastructure (PKI) services for digital certificate support, IPSecurity, and cryptography. Additionally, support is added for Transport Layer Security (TLS) that is designed to be application transparent. These functions can help provide better user authentication and help you keep your data more secure in today's challenging regulatory environment.

Refer to Software Announcement 205-034 (RFA40707), dated February 15, 2005, for a description of these previously announced functions:

- NAT Traversal for IPSecurity
- Transport Layer Security (TLS) designed to be application transparent
- CICS Sockets and Application Transparent TLS
- Integrated IPSecurity

In addition, z/OS V1.7 provides:

- **TN3270E SSLV2 protocol default change** adds the ability to specify whether Secure Sockets Layer Version 2 (SSLV2) is to be used on TN3270E connections.
- **Security Server (RACF) improvements:** The SMF unload utility can optionally create XML documents from any security event logged by RACF and EIM in SMF. These XML documents can be rendered into alternative forms, such as Web pages, making the task of analyzing the security data by security administrators and auditors more graphical and simpler.
- **Integrated Security Services Enterprise Identity Mapping (EIM) improvements:**

    - Programs written in Java that exploit the EIM Java interfaces to map authenticated user IDs to a local z/OS identity can help improve security and auditability and reduce cost.
    - Security events detected by the EIM C/C++ and Java interfaces are logged as SMF type 83 subtype 2 records, giving security administrators and auditors the data to enable tracking the use of the user ID mapping data stored in an EIM domain and to link the activities of a remote user to z/OS security events. The RACF SMF unload utility can be used to process the raw records into a tabular form or they can be included in an XML security event document.
    - EIM interfaces no longer require the caller to be APF authorized. Instead, authorization checks are made so that a caller has the authority to access EIM configuration information.
- **ICSF 64-bit support:** Integrated Cryptographic Service Facility (ICSF), in conjunction with System SSL APAR OA08775 (PTF UA14062), provided support for cryptographic hardware to be used for 64-bit callers on zSeries processors. This support for hardware-based encryption and decryption functions, which can be called automatically from System SSL, WebSphere, and Java, is expected to offer improved performance compared to the previously used software encryption functions. This support was also available in the ICSF 64-bit Virtual Support for z/OS and z/OS.e V1R6 Web deliverable, which is being replaced by the Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7 Web deliverable planned to be available in September 2005.
- **ICSF clear keys in CKDS:** Integrated Cryptographic Service Facility (ICSF) supports storing clear keys in the cryptographic key data set (CKDS), and provides the ability to specify a user-friendly label for each key. This allows application data and files to be encrypted without a clear key value having to be entered for each call to the ICSF programming interface. This can be useful when a large number of files are to be encrypted using the same clear key.
- **PKI certificate extensions:** Additional digital signature algorithms industry standards for digital certificates:
    - Certificates fulfilled by z/OS PKI Services or via the RACF RACDCERT command, support the use of the Digital Signature Algorithm (DSA). Certificate Authority (CA) keys can be DSA keys, and certificate requests can be signed with a DSA key.
    - Enhanced certificate extensions:
    - Certificate Revocation List (CRL) Distribution Point support is provided. The CRL is a list of certificates that are to be considered not valid. CRLs provide a way to verify the status of a certificate before it would expire. In addition to the Distinguished Name format, a Universal Resource Identifier (URI) format is added. With this new format, you can create certificates that contain the location of a CRL that is specified by the URI.
    - Support for CA certificates is provided, with an Authority Revocation List (ARL) Distribution Point that can be created to check the status of the CA certificates.

- Subject Alternative Name: In addition to the existing formats, such as e-mail address and IP address, a free-form format called "other name" is added. With this format, you can specify any customized name types and their corresponding values that are included in the certificates fulfilled by PKI Services.
- A PKI Services Online Certificate Status Protocol (OCSP) responder that provides the function enables applications to check the status of the certificate issued by PKI Services dynamically. OCSP relieves the need for applications that frequently download certificate revocation lists to maintain the currency of the revocation list. The OCSP support is designed to be compliant RFC 2560, enabling applications to check the validity of a certificate potentially in a more timely fashion than cached CRLs provide.
- **RACF PassTicket extensions:** An extension to SAF callable services, this support enables the generation and evaluation of RACF PassTickets. RACF PassTickets are dynamically generated password substitutes that may be used instead of a RACF password for user authentication. This extension can be invoked from C language applications and problem state programs in both 31-bit and 64-bit addressing modes without using APF-authorized services. This can help to make it easier to develop applications that may span multiple systems by providing a service on z/OS that can be utilized by the application to obtain and evaluate a RACF PassTicket for user authentication.
- **Improved security for cryptographic services:** Improvements to FTP, RACF, and z/OS UNIX are designed to provide better security for cryptographic material used by FTP in establishing SSL protected sessions. A copy of the FTP daemon's ACEE (z/OS security context) is delegated to the address space used to manage the connection to the FTP client. The FTP daemon's identity is then used in the access control decision to authorize the access to cryptographic material used in the SSL session and SSL session renegotiation. This can allow secure FTP clients to access cryptographic services without having to RACF permit each client/user to the cryptographic resource.
- **RACF FACILITY class profile for console( ):** z/OS UNIX allows users that are authorized to the new BPX.CONSOLE profile in the RACF FACILITY class to use authorized options of the _console() services (BPX1CCS and BPX4CCS) without having superuser authority. This allows you to further restrict the use of UID(0) and access to the BPX.SUPERUSER profile while allowing the use of these _console() functions in a more granular and controlled fashion.
- **Advanced Encryption Standard (AES) support for TLS/SSL applications:** The Advanced Encryption Standard (AES), proclaimed by the National Institute of Standards and Technology as the replacement for DES, is supported by the following Communications Server applications that use Transport Layer Security (TLCS) or Secure Socket Layer (SSL):
    - TN3270 Server
    - FTP client and server

FTP supports RACF-delegated profiles for protecting cryptographic resources: The z/OS FTP server supports RACF-delegated general resource profiles defined for the CSFSERV and CSFKEYS classes to protect ICSF cryptographic hardware keys and resources. By designating these resource profiles as delegated, FTP users can log in using TLS security to your z/OS FTP server without permitting every FTP client user ID to sensitive ICSF general resource profiles.

PKI Services for z/OS V1 R5 has been certified Identrus™ compliant for CA software at the Identrus 3.1 specification level. Identrus has established an open standards-based, interoperable infrastructure designed to allow companies to know who is accessing their data, services and systems. Identrus compliance establishes adherence to a common set of technology and operating standards, designed to promote interoperability within participating members, and provide a solid foundation for trust between financial institutions and their customers. PKI Services may represent savings to businesses that are currently relying on outside vendors to issue and manage digital certificates on their behalf.

About Identrus Identrus is a global solutions company that provides solutions designed to secure interactions where identity is critical. Identrus provides a worldwide network for trusted identity credentials based on global standards. The Identrus platform provides an infrastructure for financial institutions, governments and commercial entities to manage the risks associated with identity authentication Identrus membership consists of leading financial institutions.

The significance of Identrus Compliance for z/OS PKI Banking and financial service providers can face new security threats. Now that z/OS PKI Services has been certified as Identrus Compliant, banking and other financial service businesses can utilize z/OS PKI Services to establish a trusted Public Key Infrastructure. Implementing z/OS PKI Services enables businesses to serve as their own Certificate Authority (CA) enabling them to issue and manage the lifecycle of digital certificates on behalf of their business and in accordance with their security policy.

For more information, see: http://www.ibm.com/servers/eserver/zseries/security/identrus/

Slide 66

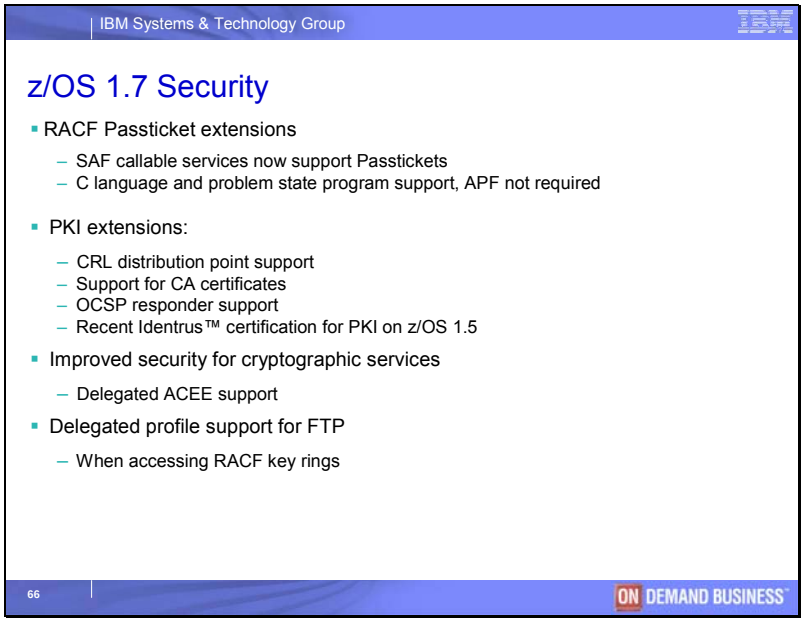

Slide 67

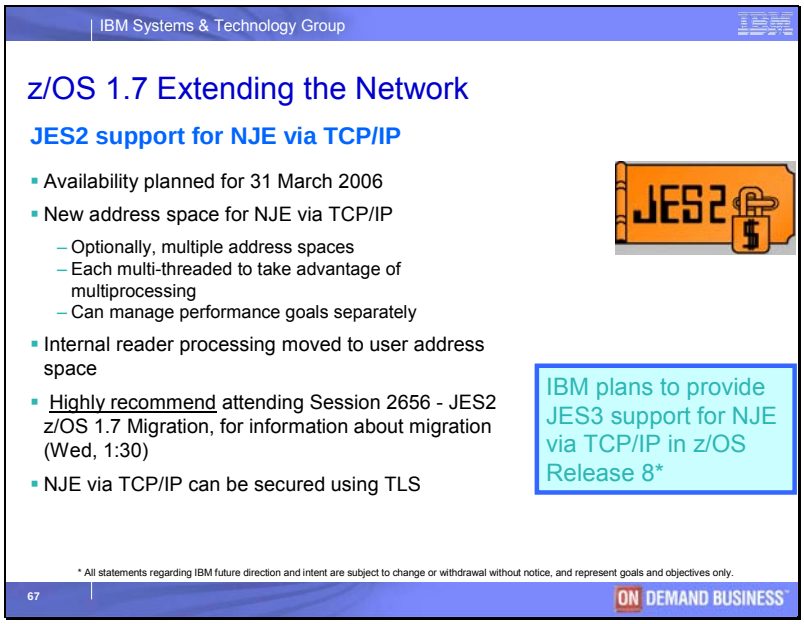

JES2 support for NJE via TCP/IP: In z/OS V1.7, JES2 supports NJE connections using TCP/IP. This can reduce the need for SNA network connections, and can also help reduce cost. Additionally, JES2 and NJE processing is now performed in separate address spaces. This can help simplify performance management and improve availability. This function is was made available in first quarter 2006 with APAR OA12364. IBM plans to provide JES3 support for NJE via TCP/IP in a future z/OS release.

Slide 68

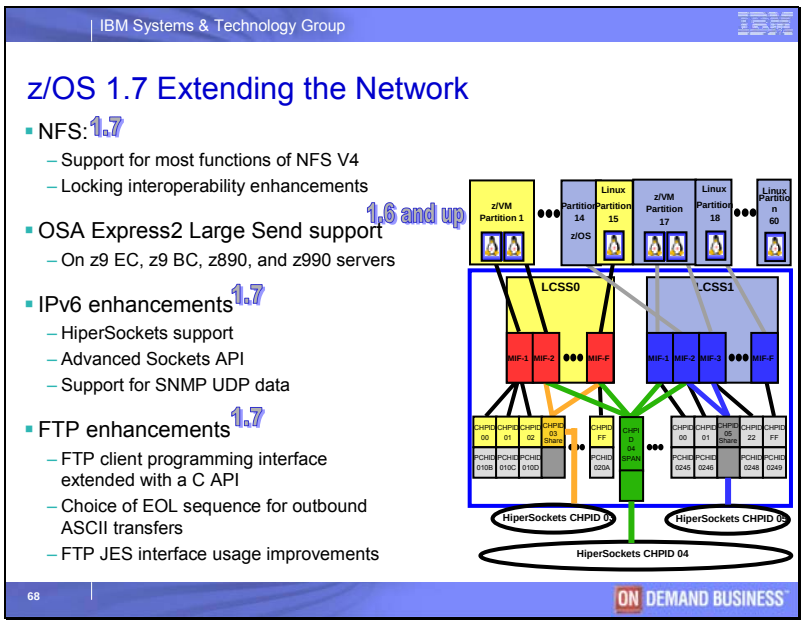

z/OS Communications Server, UNIX System Services, NFS, and JES continue to remove obstacles to growth and provide support for open standards. For example, Internet Protocol Version 6 (IPv6) is a new Internet standard that will overcome the shortage of IP addresses by increasing the IP address from 4 bytes to 16 bytes. The network growth caused by extensions to the Internet, particularly in China and other countries in the Asia and Pacific regions, combined with emerging wireless markets, will soon cause the 4-byte IP address space to become exhausted. The IPv6 standard addresses this by vastly increasing the number of available IP addresses.

z/OS Communications Server delivered the first part of IPv6 support in Release 4. Later releases added support for sockets in z/OS UNIX and in Language Environment. The final planned components of this support are included in z/OS Release 7 to support hosting applications using IPv6 in z/OS Sysplex environments.

z/OS Release 7 supports the **advanced socket APIs** described by RFCs 3542 and 2292 in C/C++, z/OS UNIX assembler callable services, and the inet6_opt_xxx() (options) and inet6_rth_xxx() (routing headers) families of functions, as well as NFS support.

FTP confidence of success level reporting: For certain types of transfers, the z/OS FTP client and server can be configured to perform additional checks and report a level of confidence that transfers have completed successfully. This is designed to provide an additional safeguard against data loss by including checks not provided for in the FTP protocol.

Slide 69

z/OS 1.7 Self-Managing Capabilities

Sysplex Distributor enhancement

- Now designed to work with WLM to direct traffic to servers that are meeting response time goals

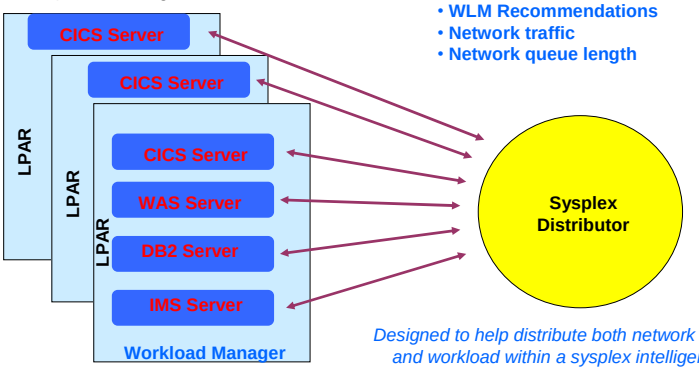

Designed to help distribute both network traffic and workload within a sysplex intelligently

Slide 70

z/OS 1.7 Extending the Network

- z/OS Load Balancing Advisor
  - Designed to interact with network-based load balancers
- Optimized Sysplex Routing
  - Better use of XCF links for messaging than in prior releases
  - Optimal route chosen for distributed traffic
- Designed to simplify configuration of TCP/IP in a sysplex
  - Improved operator commands
- FTP data transfer confidence checking and reporting
- C/C++ support for FTP client API
- CICS Sockets enhancements
  - Designed for better performance, usability, and security

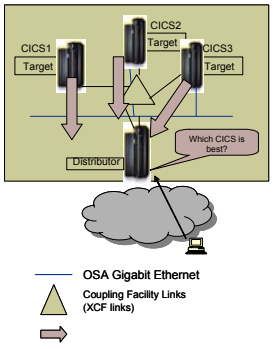

Slide 71

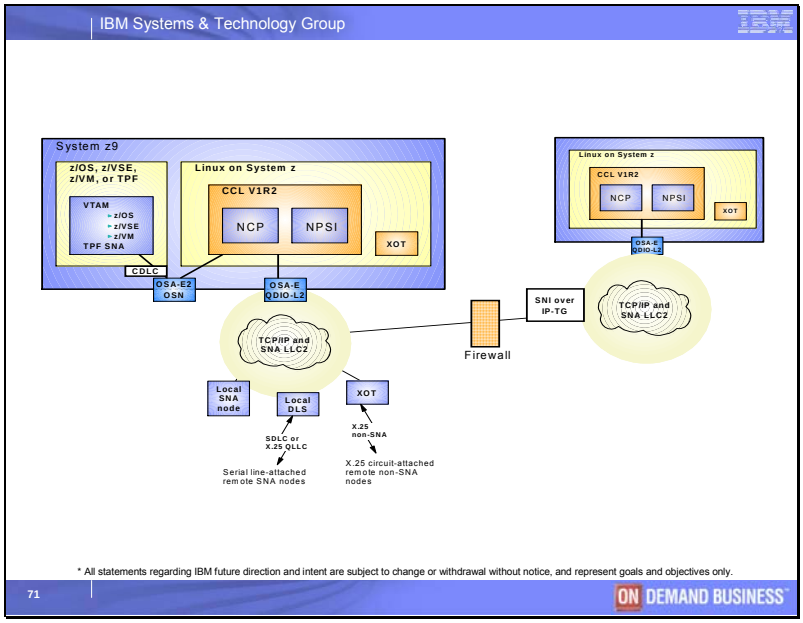

On February 15, 2005, IBM announced IBM Communication Controller for Linux (TM) on zSeries V1.1. This product is intended to provide a migration path for customers using SNA applications to communicate with business partners. In the next release of IBM Communication Controller for Linux on zSeries, IBM intends to provide enhancements in network connectivity such as channel data link control (CDLC) using OSA-Express2 OSN (OSA for NCP), data-link switching (DLSw), and an open interface for X.25 (NPSI). For more information, refer to:

http://www.ibm.com/software/network/ccl

Slide 72

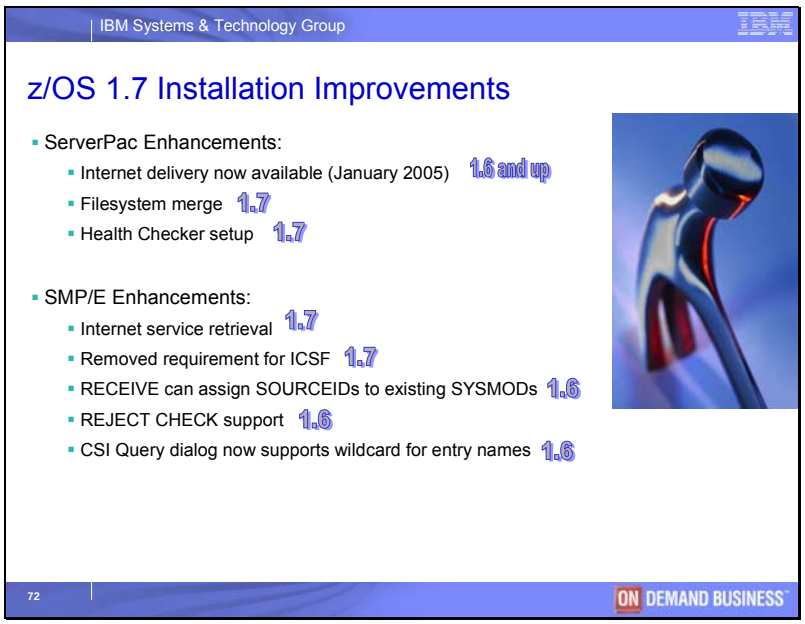

ServerPac enhancements: Internet delivery for z/OS is now available using ServerPac. (Products other than z/OS continue to be available via the Internet using CBPDO.) You can choose between tape and Internet delivery. You can also request that all components of your order be delivered electronically. You can download ServerPac orders directly to your z/OS host or use an intermediate node (such as a workstation) to transfer your order to your z/OS system.

ServerPac has also been enhanced to include:

* z/OS UNIX file system merge
* Setup for Health Checker for z/OS

Internet delivery is available for orders placed through ShopzSeries at

http://www14.software.ibm.com/Webapp/ShopzSeries/ShopzSeries.jsp

ShopzSeries is available in the United States and many other countries.

There are additional driving system requirements for Internet delivery. They are described in z/OS and z/OS.e Planning for Installation. Also, a new book, ServerPac 1.1 Planning for Internet Delivery, has been written to help provide information about setting up this new function.

SMP/E enhancements: z/OS 1.7 SMP/E includes the same functions as SMP/E V3.4. SMP/E V3.4 (5655-G44) has been enhanced with:

SMP/E Internet Service Retrieval

SMP/E Internet Service Retrieval is the latest enhancement in IBM's service delivery evolution for the z/OS platform. SMP/E V3.4 can help address the common inhibitors to Internet delivery of service, and can simplify and automate the entire z/OS service acquisition process.

SMP/E Internet Service Retrieval is designed to order and retrieve z/OS platform service from a dedicated IBM server via the Internet in one simple step right from z/OS. With SMP/E V3.4, you are now able to order and retrieve service on demand, or you can automate the service delivery process. By scheduling an SMP/E job to run once a week, or even every night, you can order and download the latest HOLDDATA and PTFs and have these services updates available when you want them.

You can request corrective service, three forms of preventive service, and even HOLDDATA.

Corrective

- You can request one or more **PTFs** by name, or
- You can request PTFs to resolve specific **APARs** . When you specify one or more APARs, you receive PTFs that resolve the APARs, but only those PTFs that are applicable to your environment.
- **Preventive** There are three selection options for preventive service:
- **Critical** service includes PTFs that resolve HIPER problems (high impact pervasive) and PTFs in error (PE).
- **Recommended** service includes all PTFs with a Recommended Service Update SOURCEID (RSUyynn) and all PTFs that resolve HIPER problems or PTFs in error (critical service).
- **All** service includes all currently available PTFs.
- **HOLDDATA** When you request **HOLDDATA** , you get only the latest two years of Enhanced HOLDDATA for the entire z/OS platform. For further information on enhanced HOLDDATA, refer to

  http://service.software.**ibm.com**/holdata/390holddata.html

Your PTF service order is customized to your SMP/E environment using a software inventory file that identifies your installed FMIDs and the PTFs already present in your environment. This means you receive the PTFs you request that are applicable to your environment plus any requisites you don't already have on hand. Furthermore, your order will not include any PTFs you already have, resulting in a comprehensive order that can be downloaded and installed quickly.

The SMP/E RECEIVE command has been extended to support this Internet Service Retrieval capability. The RECEIVE command will use the HTTP 1.0 protocol with SSL to communicate over the Internet with the dedicated IBM server. The server accepts the PTF or HOLDDATA request from SMP/E and responds with the information necessary

to download the PTF and HOLDDATA files. SMP/E then uses the FTP protocol to download the files from the IBM server to your local z/OS system.

If you currently use ShopzSeries or Service Update Facility (SUF) to order and obtain PTFs through the Internet, you should consider using SMP/E Internet Service Retrieval instead. SMP/E can provide a simpler and more automated method to acquire z/OS service than ShopzSeries or SUF.

ICSF mitigation
The z/OS Integrated Cryptographic Services Facility (ICSF) is used by SMP/E to calculate SHA-1 hash values. These hash values are calculated for files within a GIMZIP package to verify the integrity of the data within the package.

SMP/E has been enhanced to use an alternate method to calculate SHA-1 hash values if ICSF is not available for use. Although ICSF is the preferred method, SMP/E will no longer require it for use by the GIMZIP and GIMUNZIP service routines, nor for the RECEIVE FROMNETWORK or RECEIVE FROMNTS commands. If SMP/E detects ICSF is not available, then SMP/E will automatically use an SMP/E Java™ application class to calculate SHA-1 hash values as an alternative. To use the new SMP/E Internet Service Retrieval capability, SMP/E requires the IBM Software Developer Kit for z/OS, Java 2 Technology Edition, Version 1 Release 4 (5644-I56) or its logical successor. Java is not required for the installation of SMP/E, but is required at run time to use this SMP/E function.

Improved load module build
The load module build phase of the APPLY, RESTORE, and LINK LMODS commands is being enhanced to be more tolerant of allocation errors for the distribution libraries. This accommodates distribution libraries which may be offline. In this case, SMP/E continues its search for a useable copy of the module instead of immediately failing because of the error allocating the distribution library.

Note: This capability has already been provided by APAR IR53963 for existing SMP/E releases V3.2 and V3.3.

Also, in z/OS Release 6 and SMP/E V3.3:

* RECEIVE command support for assigning SOURCEID values to already-received SYSMODs
* RECEIVE FROMNETWORK support of SOCKS firewalls and secure FTP operations, using the z/OS Communications Server FTP Client
* Automatic use of IEBCOPY's COPYMOD function for load modules
* CHECK operand support for the REJECT command
* Wildcard support in the CSI Query dialog
* Internet packaging enhancements in GIMZIP and GIMUNZIP, which now support VSAM and non-VSAM data sets as well as UNIX files and directories, and to allow archives to be extracted into existing data sets

* A new service routine, GIMGTPKG, to transfer GIMZIP packages without using the RECEIVE FROMNETWORK command.

Customized Offerings Driver: The Customized Offerings Driver (5655-M12) has been updated to support the installation of z/OS 1.7 with SMP/E V3.4, the additional driving system PTFs required to install z/OS 1.7, and also supports electronic delivery for ServerPac. The Customized Offerings Driver is now at level V2.1.1.

Slide 73

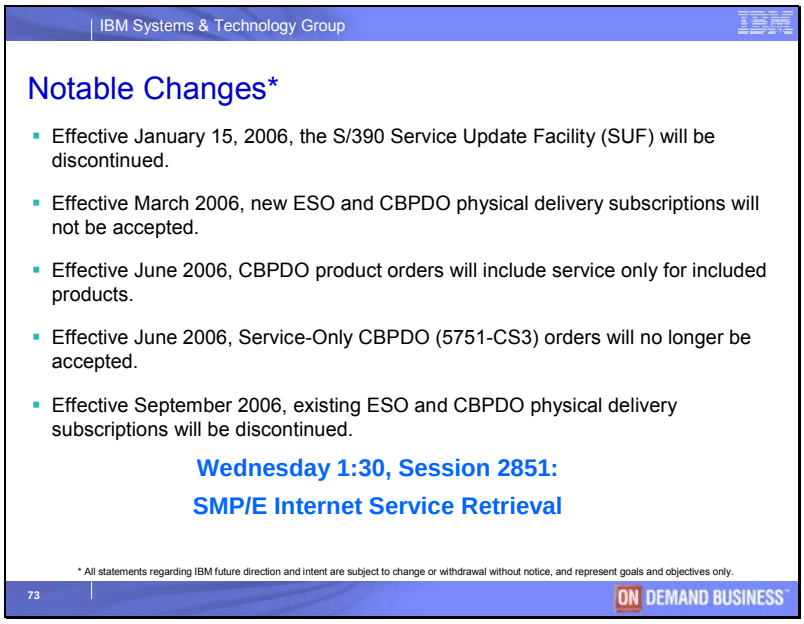

A new function in SMP/E, Internet Service Retrieval, is intended to simplify the acquisition of z/OS service. This function, which supplements existing service options in ShopzSeries, is planned to be available in September 2005. For more information about this new function, see the description of the SMP/E Internet Service Retrieval function in the section SMP/E enhancements.

Internet Service Retrieval and ShopzSeries now offer extensive options for service acquisition and delivery. To reduce the number of ordering interfaces and help assure timely delivery of new function, some older options for service delivery are being simplified or discontinued:

- Effective January 15, 2006, as previously announced in Software Announcement 205-034 (RFA40707), dated February 15, 2005, the S/390 Service Update Facility (SUF) will be discontinued.
- Effective March 2006, new ESO and CBPDO (5751-CS8 and 5775-MVS) physical delivery subscriptions will not be accepted.
- Effective June 2006, CBPDO product orders will include service only for the products included in the order. Formerly, CBPDO product orders included service

for other products licensed under the same customer number within the same SREL. To get service for other products, you can use SMP/E Internet Service Retrieval, ShopzSeries, or a fee service offering.

- Effective June 2006, Service-Only CBPDO (5751-CS3) orders will no longer be accepted. An improved option for ordering service by SREL, or for all licensed products under the same customer number, will continue to be supported in ShopzSeries. You can also get service based on what you have installed using SMP/E Internet Service Retrieval, ShopzSeries, or selected fee offerings. Note that CBPDO product orders are not affected by this change.
- Effective September 2006, existing ESO and CBPDO physical delivery subscriptions will be discontinued. You can use a job scheduling system and SMP/E Internet Service Retrieval to get service at any interval you find convenient, or use the Internet delivery subscription option available in ShopzSeries.

IBM recommends that you begin to use SMP/E Internet Service Retrieval, ShopzSeries, or one of these worldwide
fee offerings before the above changes take effect:

US - SoftwareXcel, Resolve
Europe/Middle East/Africa - Enhanced Technical Support
Canada - SupportLine
Latin America - SupportLine
Asia Pacific/South – SupportLine

For more information, visit the ShopzSeries Web site:

http://www.ibm.com/software/shopzseries

Slide 74

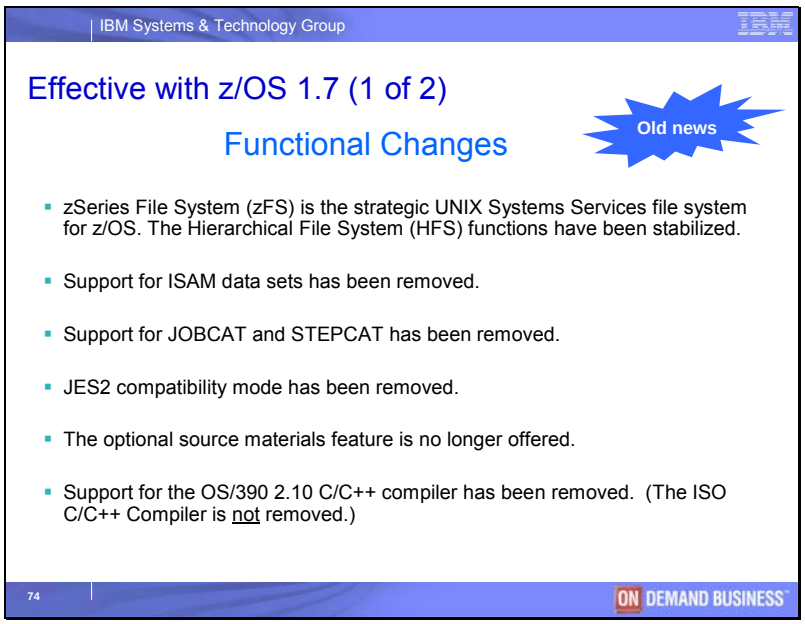

IBM plans to take the following actions effective with z/OS 1.7:

Support for ISAM data sets will be withdrawn. When this support is withdrawn, you will no longer be able to process ISAM data sets. The ISAM Compatibility Interface remains available to help you migrate to VSAM without application changes.

Support for JES2 compatibility mode will be withdrawn. As of z/OS 1.7, JES2 will no longer support compatibility with pre-z/OS 1.2 systems. The $ACTIVATE command, which converts the JES2 checkpoint from z2 mode to R4 mode (compatibility mode), will be removed. Before installing z/OS 1.7, you will need to do one of the following:

- Use the $ACTIVATE command to convert the JES2 checkpoint to z2 mode
- Offload the SPOOL, cold start JES2 z/OS 1.7, and reload the jobs on the new SPOOL

The z/OS Optional Source Code media features (which include Optional Source Code for BCP, BDT base, BDT SNA NJE, BDT File-to-File, DFSMS, MICR/OCR, BCP JPN, and Security Server RACF) will not be offered in z/OS 1.7. The last release offering these materials will be z/OS 1.6 and z/OS.e 1.6.

Support for the STEPCAT and JOBCAT JCL statements will be withdrawn. There are other facilities in DFSMSdfp (TM) allowing catalog requests to be directed to specific catalogs, and the utility of these two JCL statements has been drastically reduced by the implementation of System-Managed Storage and the placement of Unit Control Blocks (UCBs) above the 16MB line. When this support is withdrawn, any remaining JCL using these two statements will have to be changed.

The OS/390 (TM) V2.10 C/C++ compiler is planned to be removed from the C/C++ feature, leaving only the ISO C/C++ compiler. Since z/OS 1.2, the OS/390 V2.10 C/C++ compiler has been shipped in addition to the strategic ISO C/C++ compiler that is also provided with z/OS. This OS/390 V2.10 C/C++ compiler serves as a migration aid to the newest ISO C/C++ compiler.
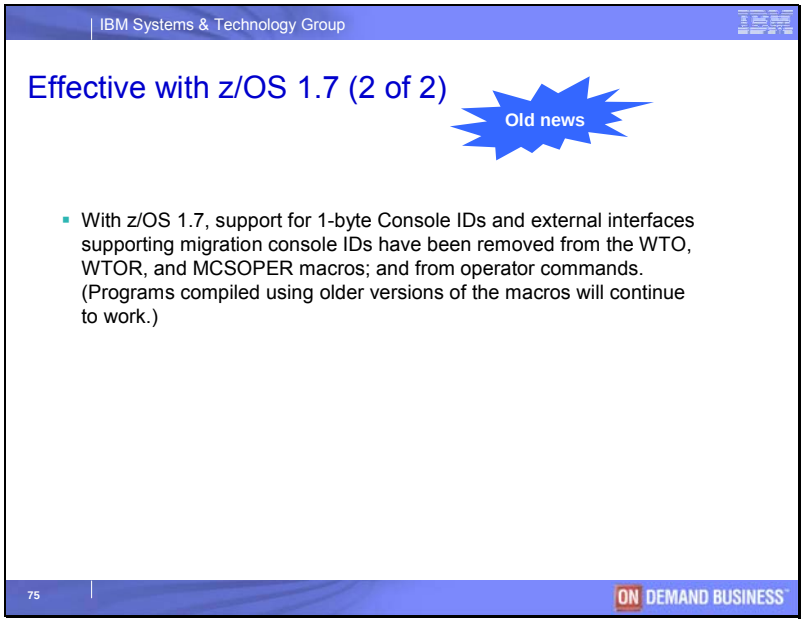
Slide 75