



International Technical Support Organization

## IBM System z9 - I/O Growth

[ibm.com/redbooks](http://ibm.com/redbooks)



# Redbooks Workshop

IBM ITSO - International Technical Support Organization

IBM System z9 Workshop

© 2005 IBM Corporation

[ibm.com/redbooks](http://ibm.com/redbooks)

International Technical Support Organization



## Notices

This information was developed for products and services offered in the U.S.A.

Note to U.S. Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of 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.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.



IBM System z9 Workshop

© 2005 IBM Corporation

2

## I/O Growth Solutions



## Large Volume Strategy

- Road Map began in early 1990's
  - **OS/390 1.3 - WLM Managed I/O Priorities**
    - Reduces impact of multiple workloads sharing same logical volumes.
  - **OS/390 1.3 - Initial PAV and Multi-Allegiance Function**
    - Reduced IOS queue time that occur as more and more data resides on a single logical volume.
  - **OS/390 2.7 - WLM Dynamic PAV Management**
    - Autonomically assigns PAV-aliases to improve parallelism where workload requires it. Reduces the number of PAV-alias subchannels required for performance when growing volume sizes.
  - **OS/390 2.8 - Native FICON with ESS**
    - Better exploitation of I/O priority by throttling bandwidth higher for high importance work.

## Large Volume Strategy (cont.)

- **OS/390 2.10 - ESS and MVS 32K Cylinder Support**
  - Volume sizes up to 27 GB. Along with WLM Alias Tuning some work loads see 50% savings in the number of subchannels required
  
- **OS/390 2.10 - Dynamic Migration to Large Volumes**
  - Dynamic rebuild of VTOC Index and expansion of VTOC
  - OW52855 and OW53053 for CVAF
  - ICKDSF Release 17
  
- **z/OS 1.1 - DCM Feature of IRD**
  - Autonomically assign bandwidth based on workload requirements for ESCON and FICON Bridge as well as support I/O priority in CSS.

## Large Volume Strategy (cont.)

- **Megamouth Storage Virtualization ("chunking")**
  - Dynamic provisioning of storage allows data growth without needing to copy the data.
- **64K Cylinder Support**
  - Continued growth of volume sizes up to ECKD architecture limit of 54 GB for 3390 track format.
- **>64K tracks for sequential datasets**
  - New JCL keyword to allow allocation of larger sequential datasets.
- **Stand-alone Dump Exploitation for Extended Format Datasets**
  - Allows better exploitation of larger volumes by placing more data on a single device.
- **Fabric I/O Priority**
  - Manages contention for shared physical resources among different workloads.
- **DCM for native FICON CHPIDs**
  - Manages bandwidth for shared physical resources among different workloads.

# Multiple Subchannel Sets

- Subchannel sets
  - Initial implementation will provide one additional set of 64K subchannels to be used exclusive of PAV-aliases devices.
  
- MIDAWs
  - New CCW type for FICON channels to eliminate data chaining penalty incurred when exploiting extended format datasets. Up to 20% response time improvements for 4K read hit. Allows line speed on 10 Gb FICON channel for large transfers (i.e. DB2 sequential prefetch).

# z/OS Data Growth

Addressability Problem

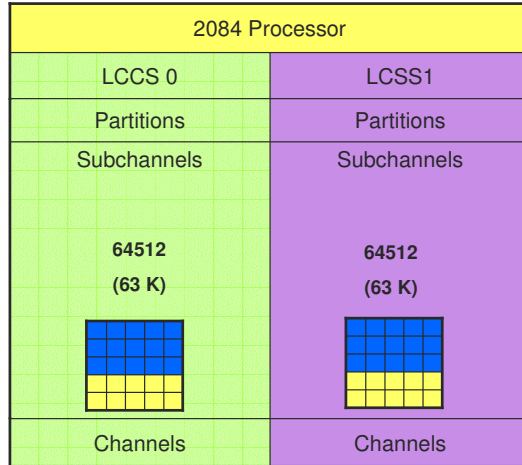
## UCBs / Subchannels

A00	A01	A02	A03	A04	A05	A06	A07	A08	A09	A0A	A0B	A0C	A0D	A0E	A0F
A10	A11	A12	A13	A14	A15	A16	A17	A18	A19	A1A	A1B	A1C	A1D	A1E	A1F
A50	A51	A52	A53	A54	A55	A56	A57	A58	A59	A5A	A5B	A5C	A5D	A5E	A5F
A90	A91	A92	A93	A94	A95	A96	A97	A98	A99	A9A	A9B	A9C	A9D	A9E	A9F
B00	B01	B02	B03	B04	B05	B06	B07	B08	B09	B0A	B0B	B0C	B0D	B0E	B0F
B40	B41	B42	B43	B44	B45	B46	B47	B48	B49	B4A	B4B	B4C	B4D	B4E	B4F
B50	B51	B52	B53	B54	B55	B56	B57	B58	B59	B5A	B5B	B5C	B5D	B5E	B5F
B60	B61	B62	B63	B64	B65	B66	B67	B68	B69	B6A	B6B	B6C	B6D	B6E	B6F




Logical Volumes


### Single Subchannel Set per LCSS (per Partition)



Subchannel per device




Base Devices




Alias Devices

### z9-109 Multiple Subchannel Sets per LCSS



■ **Set 0 – Up to 65280 (Always present)**

- Any device for allocation
- Reference in JCL
- UCB Services
- Messages
- Commands



■ **Set 1 – Up to 65535 (Optionally defined)**

- PAV alias devices only
  - Display Commands
  - Limited messages
- Designed to be compatible with existing storage CUs that support PAV



# Further Growth


## UCBs/Subchannels

Subchannel Set 0

A80	A81	A82	A83	A84	A85	A86	A87	A88	A89	A90	A91	A92	A93	A94	A95	A96	A97	A98	A99
A40	A41	A42	A43	A44	A45	A46	A47	A48	A49	A4A	A4B	A4C	A4D	A4E	A4F				
A50	A51	A52	A53	A54	A55	A56	A57	A58	A59	A5A	A5B	A5C	A5D	A5E	A5F				
A80	A81	A82	A83	A84	A85	A86	A87	A88	A89	A8A	A8B	A8C	A8D	A8E	A8F				
B80	B81	B82	B83	B84	B85	B86	B87	B88	B89	B8A	B8B	B8C	B8D	B8E	B8F				
B40	B41	B42	B43	B44	B45	B46	B47	B48	B49	B4A	B4B	B4C	B4D	B4E	B4F				
B50	B51	B52	B53	B54	B55	B56	B57	B58	B59	B5A	B5B	B5C	B5D	B5E	B5F				
B40	B41	B42	B43	B44	B45	B46	B47	B48	B49	B4A	B4B	B4C	B4D	B4E	B4F				

Subchannel Set 1

A80	A81	A82	A83	A84	A85	A86	A87	A88	A89	A8A	A8B	A8C	A8D	A8E	A8F				
A90	A91	A92	A93	A94	A95	A96	A97	A98	A99	A9A	A9B	A9C	A9D	A9E	A9F				
B40	B41	B42	B43	B44	B45	B46	B47	B48	B49	B4A	B4B	B4C	B4D	B4E	B4F				
B50	B51	B52	B53	B54	B55	B56	B57	B58	B59	B5A	B5B	B5C	B5D	B5E	B5F				



Logical Volumes

# Super Volumes


## UCBs/Subchannels

Subchannel Set 0

A80	A81	A82	A83	A84	A85	A86	A87	A88	A89	A8A	A8B	A8C	A8D	A8E	A8F				
C00	C01	C02	C03	C04	C05	C06	C07	C08	C09	C0A	C0B	C0C	C0D	C0E	C0F				
D00	D01	D02	D03	D04	D05	D06	D07	D08	D09	D0A	D0B	D0C	D0D	D0E	D0F				

Subchannel Set 1

A80	A81	A82	A83	A84	A85	A86	A87	A88	A89	A8A	A8B	A8C	A8D	A8E	A8F				
B40	B41	B42	B43	B44	B45	B46	B47	B48	B49	B4A	B4B	B4C	B4D	B4E	B4F				



Logical Volumes

## Software support

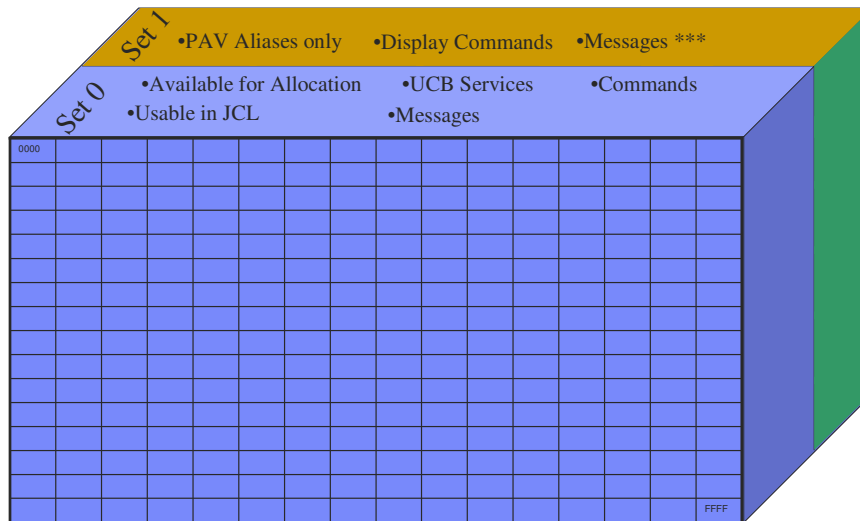
- Overview
- Multiple Sub-Channel Set definition
  - General information
  - Command changes
  - New messages
- Dependencies
- Migration & coexistence
- Installation

## Overview

- **Data growth within a sysplex is constrained by device addressability:**
  - ~64K subchannels available for use per z/OS image
  - Technology made available throughout the years reduces subchannel availability
    - Parallel Access Volumes
    - PPRC
    - GDPS and Hyperswap
    - XRC
    - Flashcopy
- **Solution:**
  - ▶ Provide additional Subchannels that can be used by the system with minimal effect on applications
- **Benefit:**
  - ▶ Device addressability constraints are relieved, minimal or no impact



## Multiple Subchannel Sets Definition



## MSS – General Information

- "Logical" 5-digit devices
  - Accepted as command inputs, inputs to GTF
  - messages and formatted outputs
  - UCBDEVN macro
- New messages for 5-digit devices
  - Messages for devices in subchannel set 0 are unchanged.
- System Trace and GTF Trace support logical 5-digit devices
- Programming interfaces have new keywords (UCBSCAN, UCBLOOK, UCBINFO, UCBDEVN)
  - Device number remains 2 bytes
    - UCBCHAN is 2 bytes and represents the device number
  - IOSDUPI and IOSDUPFX mappings contain UCBSSID
  - UCBCHAN in combination with UCBSSID logically represent the device number
  - Only PAV Aliases can have non-zero UCBSSID values
    - May allow other devices in future releases
  - SMF and ENF data is different for PAV aliases in subchannel set 1
- Only ECMB mode is supported on z/OS 1.7 (no ECMB=YES|NO parameter in IEAOPTxx)

## Command Changes

- D M=DEV output shows information on devices in both subchannel sets
  - Input device numbers can be 1-5 hex digits
  - Output matrix enlarged to show devices in both subchannel sets
- DEVSERV enhanced to allow specification of devices in both subchannel sets
  - Output data includes aliases in both subchannel sets

## New “Buddy” Messages

IEA1522I    IOS1120D  
 IEA1525W    IOS1124D  
 IEA1774I    IOS1150I  
 IOS1000I    IOS1151I  
 IOS1001E    IOS1152E  
 IOS1002A    IOS1153E  
 IOS1003A    IOS1154I  
 IOS1017I    IOS1155I  
 IOS1050I    IOS1156I  
 IOS1051I    IOS1157I  
 IOS1070E    IOS1158I  
 IOS1071I    IOS1159I  
 IOS1075E    IOS1164I  
 IOS1076E    IOS1202E  
 IOS1077E    IOS1203I  
 IOS1078I    IOS1207I  
 IOS1079I    IOS1208I  
 IOS1080I    IOS1250I  
 IOS1109E    IOS1251I  
 IOS1110D    IOS1291I  
 IOS1111D    IOS1444I  
 IOS1112D    IOS1500I  
 IOS1116A    IOS1502I  
 IOS1117A    IOS1515A  
 IOS1118A    IOS1522A  
 IOS1119A


IEA522I UIM CBDUS022 DID NOT BUILD A DFT FOR DEVICE NUMBER 000B  
 IEA1522I UIM CBDUS012 DID NOT BUILD A DFT FOR DEVICE NUMBER 1001C  
  
 IOS164I DEVICE 0999, NO MATCHING UCB FOUND  
 IOS1164I DEVICE 10999, NO MATCHING UCB FOUND

## Interactions & Dependencies

- Hardware
  - ▶ MSS is a new function on the IBM System z9
- Software
  - ▶ HCD allows migration of PAV aliases into subchannel set 1
  - ▶ Applications will see no difference. The system will exploit this function if enabled.
  - ▶ See HCD/HCM presentation for more information on alias definition in subchannel set 1

## Migration/Coexistence Considerations

- z/OS 1.7 and the IBM System z9 allow subchannel set exploitation
- Migration of aliases will require updates to IODF
  - Once aliases are moved, duplicate 4-digit device numbers become possible



## Installation

- Prerequisites for installation:
  - ▶ Install OA08197 on lower levels of z/OS (not specific to MSS but required for IODF changes).
  - ▶ Install OA07875 on lower levels (HCD).
  
- Publications References:
  - ▶ None

## Appendix

- Important Publication Changes for IOS MSS Support:
  - SA22-7627-12 z/OS V1R7.0 MVS System Commands
  - SA22-7636-09 z/OS V1R7.0 MVS System Messages, Vol 6 (GOS-IEA)
  - SA22-7638-10 z/OS V1R7.0 MVS System Messages, Vol 8 (IEF-IGD)
  - SA22-7610-09 z/OS V1R7.0 MVS Authorized Assembler Services Reference ENF-IXG
  - SA22-7612-06 z/OS V1R7.0 MVS Authorized Assembler Services Reference SET-WTO
  - SA22-7607-08 MVS Assembler Services Reference IAR-XCT