

 #SHAREorg

What We've Done for You Most Recently with PDSE

Cecilia Carranza Lewis
carranc@us.ibm.com
IBM

February 5, 2013
Session 12979





Disclaimer

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.



Agenda

- Quality Improvements and Trends
- Recent Enhancements
- What's New in z/OS V2.1
- **Don't Miss**
 - **Session 12981: Using PDSEs in Your SYSPLEX: Best Practices and Troubleshooting [Friday 9:30AM]**
 - **Session 12982: PDSE Performance Basics [Friday 11AM]**

Quality Activities – Historical Issues of the mid-90's

- **Issues**
 - Outages due to excessive use of Extended Common Storage Area (ECSA)
 - System hangs due to failure and cancel situations
 - No tools to display/diagnose problem situations
- **Solutions**
 - PDSE Address Space (SMSPDSE)
 - SPE delivered in 2002 On OS/390 R10 thru z/OS V1R5
 - *Over 98% of PDSE structures moved from Common into address space*
 - Improved Diagnosis tooling
 - PDSE Monitor tool: Monitors the PDSE address space
 - Analysis commands: Detects a number of the most common latch and lock contention problems
 - Freelatch command: Releases any latch that the ANALYSIS command indicates is held
 - Enhanced internal test
 - Developed internal test tool, Valkyrie
 - Rebuilt and fortified PDSE test workloads
 - *Including, addition of robust cancel scenarios*
- **Result**
 - Eliminated outages due to excessive ECSA usage (increased scalability)
 - Reduced IPLs due to system hangs
 - Provided storage administrators tools for monitoring and diagnosing
 - Increased scalability

Quality Activities continued

- Continued focus on high availability and outage eliminations
- **Solution**
 - PDSE Restartable Address Space (z/OS V1R6)
 - Non-disruptive, transparent restart
 - New address space (SMSPDSE1)
 - Enable via Parmlib member IGDSMSxx
 - PDSESHARING(EXTENDED) and PDSE_RESTARTABLE_AS(YES)
 - Restart
 - Vary SMS,PDSE1,Restart
- **Result**
 - Further reduced outages and need to re-IPL

Continued focus on high availability, outage eliminations, and scalability



- **64-bit Virtual Storage**
 - Directory buffers and control blocks for member connections moved to 64-bit virtual, allowing for greater number of concurrently opened members as well as scalability for caching directory buffers.
- **Buffer Beyond Close**
 - Provides an option to retain buffers (directory and member data) after the last close of a PDSE dataset.
 - Improves performance for PDSEs that are opened/closed repeatedly.
- **Error Redrive**
 - On an ABEND0F4 due to invalid directory cache, enhance PDSE error recovery to percolate back to the entry of PDSE, purge the cache, and redrive the PDSE request....thus eliminating the 0F4.
- **Identifying corrupted PDSE's in LNKLIST**
 - Validates PDSEs in LNKLIST using internal validation tooling; Issues message (IGW037I) identifying any corrupted PDSE data sets at IPL time.

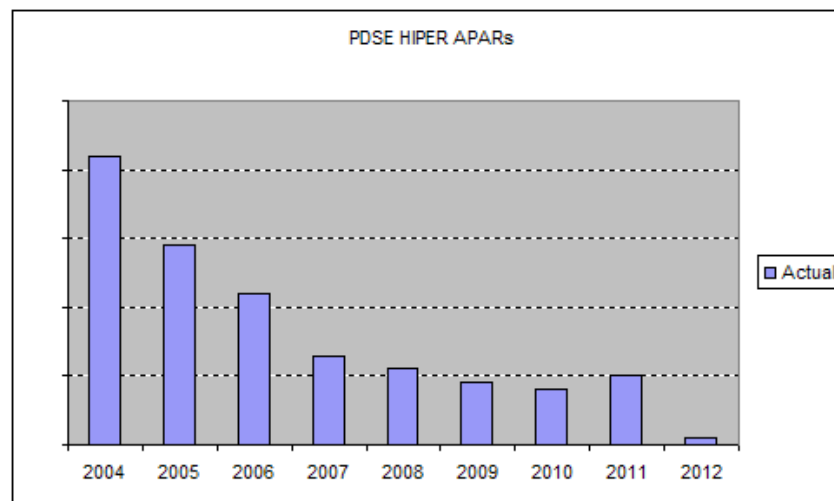
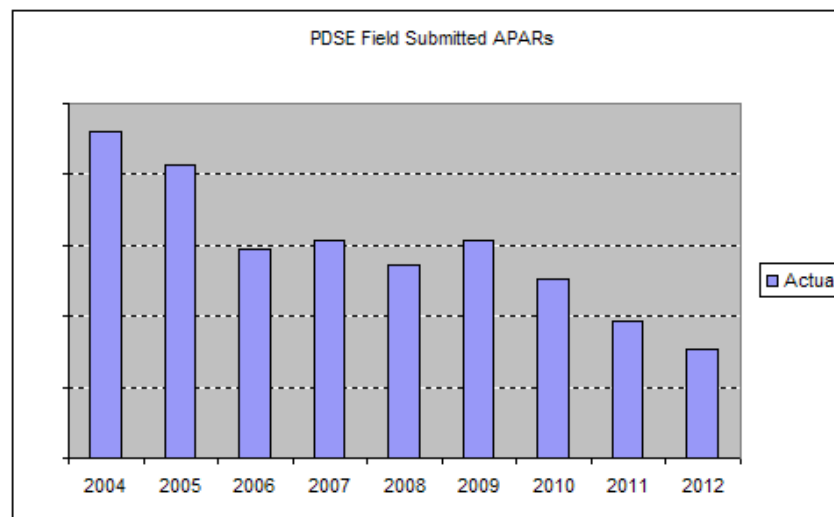
Continued focus on high availability, outage eliminations, and scalability cont.



- **Monitor/Analysis Command Improvements**
 - On Monitor, new DISPLAY and DUMP options
 - DISPLAY turns monitor on and causes any error messages displayed on console; DUMP turns monitor on and causes dump to be taken on next error
 - Analysis enhanced to provide option to take a dump if any problem found
- **PDSE DISPLAY SMS Enhancement**
 - Enhanced DISPLAY SMS command with new options to display point-in-time PSDE cache information in real time and overall effectiveness of PDSE caching
- **Data Set Level Caching Statistics**
 - Provided data set caching statistics in SMF Type 14 and 15 records
 - A new optional section called *PDSE Data Set Caching Statistics (type6) for Directory and Members*

PDSE Quality Trends

- Results of PDSE improvements
 - Field APAR rate has been reduced by nearly **67%** since 2004
 - Significant reduction of **98%** in HIPER APARS
- In addition,
 - PMR rate has been reduced by **62%** since 2004



Recent Enhancements prior to V2.1

Enhancements: z/OS V1R13 IEBPDSE

- **IEBPDSE Utility**
 - PDSE Validation tool designed to verify the structure of the PDSE directory.
 - First phase of the validation tool was integrated into z/OS 1.12 to identify corrupt PDSE's in LNKLST and at NIP time.
 - Invoked via JCL: EXEC PGM=IEBPDSE
 - Optional PARM keyword: PARM [DUMP|NODUMP]
 - *DUMP option indicates that the PDSE validation utility is to issue an ABEND in the PDSE address space when an error has been found in the analysis of the PDSE.*

Enhancements: z/OS V1R13 Commands

- **Requirement**
 - Customers expressed the need to discard cached directory pages in order to recovery from possible incorrect pages in the directory cache.
- **Solution**
 - Provide a command to discard all buffered data for a PDSE
 - V SMS,PDSE(1),REFRESH
 - Messages
 - IGW052I - REFRESH is successful
 - IGW053I - REFRESH is unsuccessful

Enhancements: z/OS V1R13 Commands

- **Requirement**
 - Customers expressed the need to determine the users of any given PDSE, and determine which users are blocking data set access.
- **Solution**
 - Provide a command to display the connections for a PDSE. Used to help determine scope of a problem with a particular PDSE. Allows Level 2 and customer to decide whether to cancel jobs or restart PDSE address space.
 - Display
SMS,PDSE(1),CONNECTIONS,DSNAME(dsn),VOLSER(volume)
 - Message
 - IGW051I PDSE CONNECTIONS REPORT

Enhancements: z/OS V1R13 Fairness

- **Problem**
 - Today, a writer to a PDSE can be blocked by high read activity against the PDSE on other systems.
- **Solution**
 - Fairness is designed to improve the efficiency of sharing a PDSE when it is being shared across systems and updated repeatedly. The PDSE locking structure is changed to prevent one system from dominating access to a PDSE that is being accessed across systems.
- **Benefit**
 - Performance of PDSE is more consistent
 - There may be less GRS activity associated with PDSE processing

What's New in z/OS V2.1

V2.1 Easing Member Size Limitation

- PDSE Larger Member Size
 - Currently PDSE members are limited to 15,728,639 records.
 - This limit does not exist for PDS datasets.
 - **New enhancement:** Increase the limit on PDSE member size.
 - PDSE member size is planned to be over 125 times larger (approximately 2,146,435,071 records) than the current limit in many circumstances, and substantially larger than the maximum supported size of a PDS member.
 - **SHARE Requirement met: SSMVSS11010**
- **? Why it Matters:** Provide additional scalability and usability benefits of using PDSEs in place of PDSs and make it feasible to use PDSEs instead of multiple large sequential data sets.

Accessing Records Beyond Current Limit with BSAM/BPAM

- A program can WRITE records beyond current limit under the following conditions:
 - DSORG=PS
 - DCB MACRF is W (not WP) <WP: Write with Note/Point>
 - Or
 - DCB MACRF is WP and DCBE BLOCKTOKENSIZE=LARGE*
 - DSORG=PO
 - DCBE BLOCKTOKENSIZE=LARGE*

(*) If BLOCKTOKENSIZE=LARGE not specified, a WRITE beyond record 15,728,639 will result in ABEND S002-A8.

Note: BLOCKTOKENSIZE=LARGE indicates that program can support a 32-bit Note token.

Accessing Records Beyond Current Limit with BSAM/BPAM cont.

- A program can READ records beyond current limit in all cases (regardless of DSORG or MACRF)
- A program can NOTE records beyond the current limit under the following conditions:
 - DCBE BLOCKTOKENSIZE=LARGE*

(*)If BLOCKTOKENSIZE=LARGE not specified, a NOTE issued for a record after record 15,728,639 will result in ABEND 002-A8.

Accessing Records Beyond Current Limit with QSAM



- Any QSAM program can get/put records beyond the current limit in all cases.

Coexistence with Down-Level Systems (V1R12 and V1R13)

- Issuing a WRITE/PUT beyond record 15,728,639 will result in existing ABEND S002-A8.
- Issuing a READ/GET beyond record 15,728,639 will be allowed.

Coexistence with Down-Level Systems cont.

- Issuing a NOTE beyond record 15,728,639 with DCBE BLOCKTOKENSIZE=LARGE will return a valid NOTE value.
- Issuing a NOTE beyond record 15,728,639 without DCBE BLOCKTOKENSIZE=LARGE will result in ABEND S002-A4 (with coexistence apar applied, otherwise unpredictable).
- Issuing a POINT beyond record 15,728,639 with DCBE BLOCKTOKENSIZE=LARGE will result in out-of-extent on next READ/WRITE.

Messages and Codes

- IEC036I 002-A8
 - Existing description
 - “Maximum number of records in a PDSE member has been exceeded. Or maximum number of blocks allowed in an extended sequential data set has been exceeded.”
 - Additional description added
 - “Or a Note was issued for a PDSE data member and the record is more than 15,728,639 records into the member and BLOCKTOKENSIZE=LARGE was not specified when the PDSE was opened. NOTE could not have returned a valid value.”

New NOTE/POINT Token Range

- PDSE Directory
 - X'00000001'
- PDSE Members (MLT)
 - X'00000002' to X'0007FFFFFF'
- PDSE Record Numbers (RLT)
 - X'00100001' to X'7FFFFFFFFF'

V2.1 PDSE Code Restructure

- Current issues
 - Performance degradation for specific processing compared to PDSs.
 - Excessive use of auxiliary storage utilization when PDSEs are updated
 - Result of IEBCOPY may be larger than needed
- Solution
 - Improve performance when processing large PDSEs by shortening pathlengths and enhancing directory structures, and more efficient aux storage management.

Restructure Focus

- PDSE processing restructured with goals of
 - Reducing the CPU utilization for most common processing primarily due to the way PDSE directories are processed.
 - Improvements in index updates
 - The reduction most significant for PDSE with larger number of members (>1000 members)
 - Improving use of aux storage when PDSEs are updated
 - Cache algorithm improved when dealing with updates to be more efficient in releasing unused buffers.

Customer view

- No external changes required to exploit this enhancement
- No data set structure changes
- No incompatibilities
- No coexistence considerations

- Improvements
 - CPU reduction
 - Less aux usage (for updates)
 - Less I/O due to improvement in directory processing

Performance

- Reduced CPU and Auxiliary Storage requirements for index processing.
- TSO Performance Workload executing many concurrent browse/edit operations on very large PDSE datasets. Performance results showed improvement from V1R13 to V2R1:
 - 87% CPU reduction and 87% storage reduction in SMSPDSE address space.
 - 55% CPU reduction and 43% storage reduction in TSO User address spaces.

Note: Performance improvements are based on internal IBM laboratory tests. Your results will vary.

Performance

- TSO Workload
 - 300 Concurrent Users accessing 30 large PDSEs
 - Mixture of RECFMs FB, VB, U
 - Each PDSE contained over 10,000 members
 - Varied sizes for each member
- TSO Response Time Improvements
 - Enter browse, view member list
 - 60% improvement due to improved index processing
 - Edit existing member
 - 30% improvement
 - Edit new member
 - 30% improvement

Note: Performance improvements are based on internal IBM laboratory tests. Your results will vary.

V2.1 PDSE Version 2



- PDSE Version 2 Enhancements
 - **New enhancement:** Simplify PDSEs to be able to make partial release more effective and to improve general PDSE performance.
 - The system will create a new PDSE format, Format 2, which will enhance partial release, consolidate directory pages when possible, improve read performance, and reduce virtual storage utilization for PDSE processing.
 - *Toleration of the new PDSE format is planned for z/OS V1.12 and z/OS V1.13*
 - Overall PDSE performance will be improved.
 - *The path length of almost all PDSE related operations will be reduced, and index searches will be improved.*
 - *Unnecessary structures from the directory will be removed allowing space to be used more efficiently.*
- **? Why it Matters:** These enhancements are intended to provide additional scalability and usability benefits of using PDSEs in place of PDSs, make it feasible to use PDSEs instead of multiple large sequential data sets, and help reduce the space required for PDSEs.

Customer View

- Benefits
 - Majority of processing to benefit from this support due to improved space utilization, reduction in CPU processing and I/O.
 - First OPEN of large PDSE should see improved performance
 - For PDSE with variable records, creating large members should be much faster
- However,
 - For PDSE with variable records, when OPEN followed by ‘blind’ Point to end of member, read will be slower

How to create PDSE Version 2

- New option on DSNTYPE keyword
 - DSNTYPE=(LIBRARY,{1|2}) where
 - 1 – existing type of PDSE (default)
 - 2 – new version PDSE
 - Supported on JCL, TSO Allocate
- New option in IGDSMSxx member of SYS1.PARMLIB
 - DSNTYPE=({LIBRARY|PDS|HFS},{1|2})
- Precedence
 - DSNTYPE on JCL takes precedence over PARMLIB

How to create PDSE Version 2 cont.

- **LIKE=**
 - If DSNTYPE was specified on the JCL but version was not, the value for version will be obtained from the IGDSMSxx member NOT from the LIKE= data set.
 - If DSNTYPE was not specified on the JCL, then the DSNTYPE value and version will be obtained from the LIKE= data set.
 - If DSNTYPE was not specified on the JCL and LIKE= was not specified then the DSNTYPE value will be obtained from the DATA CLASS and, if necessary, the version number will be obtained from the IGDSMSxx member.
 - If the DSNTYPE is not available from any source (JCL, LIKE= or DATA CLASS) then the default values for both DSNTYPE and the version (if applicable) will be picked up from the IGDSMSxx member.

How to create PDSE Version 2 cont.

- REFDD=
 - Use the REFDD parameter to specify attributes for a new data set by copying attributes of a data set defined on an earlier DD statement in the same job.
 - One of the attributes which are copied to the new data set from the attributes specified on the referenced DD statement is DSNTYPE. DSN version will also be copied to the new data set from the referenced DD statement.

How to detect PDSE Version 2

- ISMF
 - Data Set List: Version added to data under column 'DATA SET NAME TYPE'
- ISITMGD
 - New field: ISMDSNVER
- SMF Type 14/15
 - New field: SMF14DSVER

How to detect PDSE Version 2 cont

- ISMF – Data Set List example
 - Version displayed with data set type

```
-----
DGTLP13                                DATA SET LIST
Command ==>                                Scroll ==> CSR
                                           Entries 22-38 of 38
                                           Data Columns 30-33 of 42
Enter Line Operators below:                **ENTRIES HIDDEN**
**FILTERED LIST**
LINE                                     DATA SET  NUM OF  ENTRY  REBLK
OPERATOR                                DATA SET NAME  NAME TYPE  STRIPES  TYPE  IND
--- (1) ---  ----- (2) -----  -- (30) ---  - (31) --  -- (32) --  (33) -
      SYS1.LINKLIB                OTHERS                --  NONVSAM  NO
      SYS1.LINKLIB.PDSE0          LIBRARY,1            --  NONVSAM  NO
      SYS1.LPALIB                OTHERS                --  NONVSAM  NO
      SYS1.LINKLIB.PDSE2          LIBRARY,2            --  NONVSAM  NO
```

Coexistence

- Allow down-level systems (z/OS V1R12 and z/OS V1R13) to access PDSE Version 2 data sets
 - However, down-level systems cannot create a new PDSE Version 2 data set

Performance

- TBD (In progress)

References for z/OS V1R13

- Partitioned Data Set Extended Usage Guide, RedBooks
 - SG24-6106-01 or Ibm.com/redbooks
- z/OS DFSMS: Using Data Sets
 - SC26-7410-10
- z/OS DFSMS Macro Instructions for Data Sets
 - SC26-7408-09
- z/OS DFSMSdfp Utilities
 - SC26-7414-07
- z/OS MVS System Management Facilities (SMF)
 - SA22-7630-22
- Why You Want to Use Partitioned Data Set Extended (PDSE)
 - Techdoc PRS4280

System z Social Media Channels

- Top Facebook pages related to System z:
 - [IBM System z](#)
 - [IBM Academic Initiative System z](#)
 - [IBM Master the Mainframe Contest](#)
 - [IBM Destination z](#)
 - [Millennial Mainframer](#)
 - [IBM Smarter Computing](#)
- Top LinkedIn groups related to System z:
 - [System z Advocates](#)
 - [SAP on System z](#)
 - [IBM Mainframe- Unofficial Group](#)
 - [IBM System z Events](#)
 - [Mainframe Experts Network](#)
 - [System z Linux](#)
 - [Enterprise Systems](#)
 - [Mainframe Security Gurus](#)
- Twitter profiles related to System z:
 - [IBM System z](#)
 - [IBM System z Events](#)
 - [IBM DB2 on System z](#)
 - [Millennial Mainframer](#)
 - [Destination z](#)
 - [IBM Smarter Computing](#)
- YouTube accounts related to System z:
 - [IBM System z](#)
 - [Destination z](#)
- Top System z blogs to check out:
 - [Mainframe Insights](#)
 - [Smarter Computing](#)
 - [Millennial Mainframer](#)
 - [Mainframe & Hybrid Computing](#)
 - [The Mainframe Blog](#)
 - [Mainframe Watch Belgium](#)
 - [Mainframe Update](#)
 - [Enterprise Systems Media Blog](#)
 - [Dancing Dinosaur](#)
 - [DB2 for z/OS](#)
 - [IBM Destination z](#)
 - [DB2utor](#)





SHARE
Technology • Careers • Skills

THANK YOU

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



Legal Disclaimer



NOTICES AND DISCLAIMERS

(C) 2012, 2013 IBM Corporation

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products in connection with this publication 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.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



Trademarks



The following are trademarks of the *International Business Machines Corporation*:

IBM, DFSMS/MVS, DFSMSHsm, DFSMSrmm, DFSMSdss, DFSMSopt, DFSMS Optimizer, z/OS, eServer, zSeries, MVS, FlashCopy®

The information contained in this presentation is distributed on an 'AS IS' basis without any warranty either expressed or implied, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. The use of this information is a customer responsibility and depends on the customer's ability to evaluate and integrate it into the customer's operational environment.

 #SHAREorg



What We've Done for You Most Recently with PDSE

Cecilia Carranza Lewis
IBM

February 5, 2013
Session 12979

