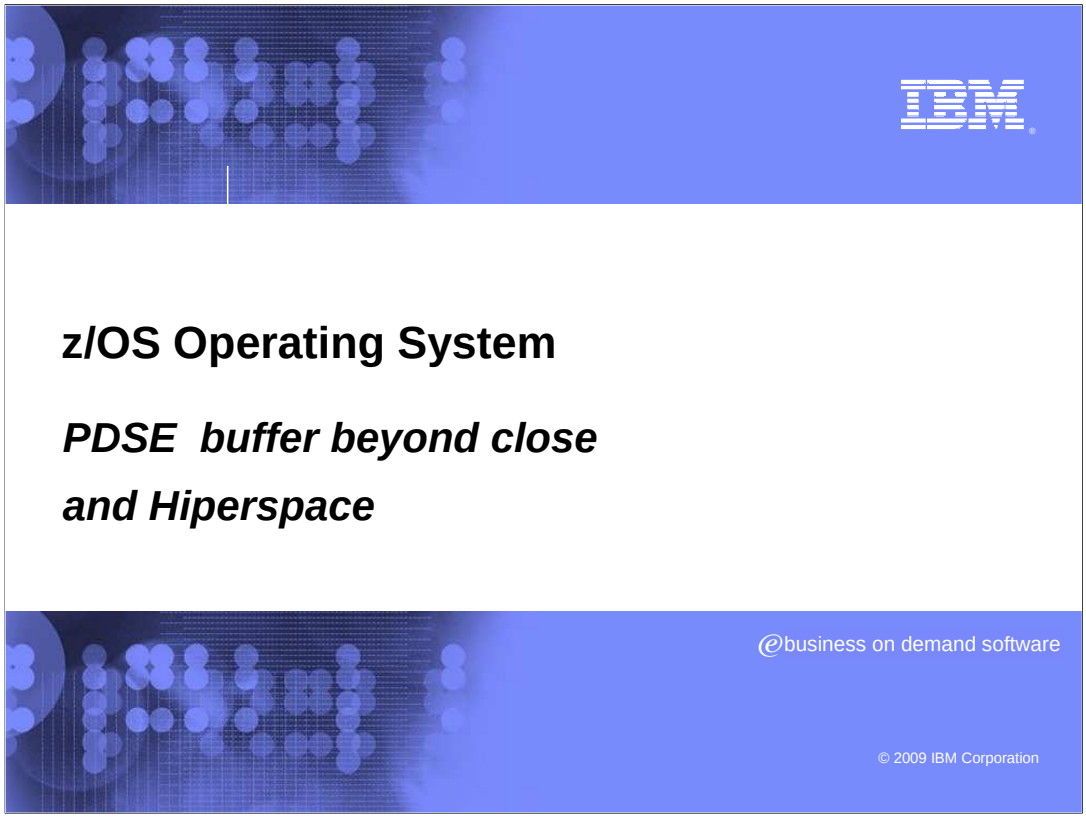

This education topic provides information regarding PDSE Buffer Beyond Close and Hiperspace™.

Agenda

- Problem / solution
- Usage and invocation
- Performance considerations
- Validation
- PDSE cache statistics
- LLA / VLF member caching
- References

This presentation will begin with a discussion of the problem in PDSEs and the solution that PDSE buffer beyond close brings; Then it covers how to use this new feature; afterwards it will talk about the performance considerations and they to validate this new feature (using SMF records). It will then move on to discuss LLA/VLF Member caching and what has been done to improve the performance of processing PDSE members. Finally, a list of References is provided.

PDSE buffer beyond close

- Problem:
  - When a PDSE is no longer open on a system the buffers for the PDSE are purged from in-storage cache.
  - PDSEs that are opened and closed repeatedly do not get as much benefit from caching buffers as they could.
- Solution:
  - Provide an option to retain buffers after the last close of a PDSE.
- Benefit:
  - Improved performance if the PDSE is opened within 15 minutes after the last close.

Read problem in the slide.

The buffers for a PDSE are purged from in-storage cache when a PDSE is closed (the last close on a system).

Read Solution in the slide.

With the option of PDSE Buffer Beyond Close, you can retain the buffers after the last close of a PDSE dataset. This option will eliminate the overhead of an I/O operation because the buffers that you are trying to access remain in cache.

Read Benefit in the slide.

This will improve the performance of a PDSE that is opened soon after the last close.

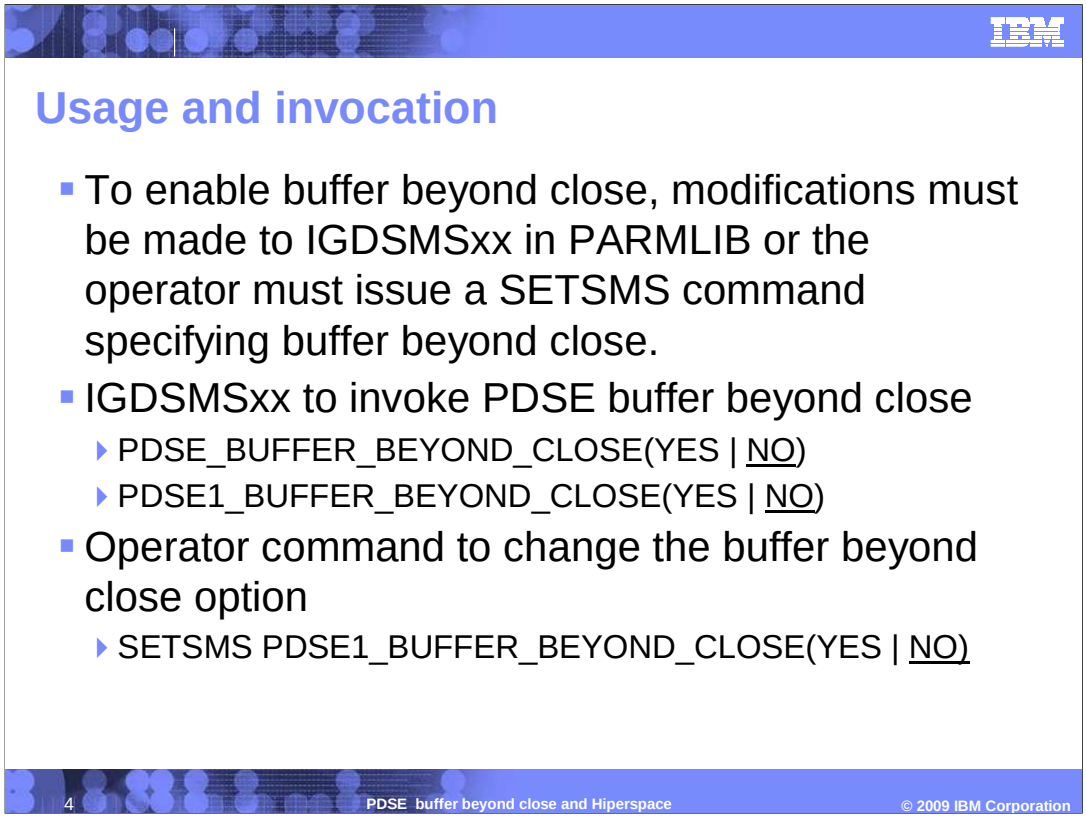

The following are the parameters that help you to set SMSPDSE and SMSPDSE1 address spaces with the ability of buffer beyond close.

PDSE_BUFFER_BEYOND_CLOSE (YES | NO) or

PDSE1_BUFFER_BEYOND_CLOSE(YES| NO)

For the (SMSPDSE | SMSPDSE1) address space, this parameter specifies whether to keep directory and member data in storage beyond the last close on this system of a PDSE dataset. If the NO option is selected or defaulted a PDSE's directory and member data will be purged from the in-memory cache when the last close of the data set occurs. If the YES option is selected, then the PDSE's directory and member data will be retained in the in memory cache beyond the last close of the data set.

There are two ways of enabling Buffer Beyond Close:

1. Modify IGDSMSxx in PARMLIB to invoke PDSE buffer beyond close. The fields to be changed are:

   PDSE_BUFFER_BEYOND_CLOSE(YES | NO) : No is the default. This will apply to the SMSPDSE Address Space. If you have only one PDSE AS, changing it to Yes will help to decrease the amount of I/O. If there are two Address Spaces, SMSPDSE will be only for PDSEs in LNKLIST and there will be no benefit for buffer beyond close.

   PDSE1_BUFFER_BEYOND_CLOSE(YES | NO): No is the default. This will apply to SMSPDSE1 Address Space. To benefit from buffer beyond close, this field needs to be changed to Yes.

2. SETSMS PDSE1_BUFFER_BEYOND_CLOSE(YES | NO): This command is issued from the console and will take effect immediately.

Performance considerations

- Retaining PDSE buffers longer could raise a virtual storage concern.
- PDSE Cache Least Recently Used (LRU) Tuning Knobs.
  - SMSPDSE Address Space
    - PDSE_LRUTIME
    - PDSE_LRUCYCLES
  - SMSPDSE1 Address Space
    - PDSE1_LRUTIME
    - PDSE1_LRUCYCLES

Be advised that retaining PDSE buffers longer could raise a virtual storage concern. Without Buffer beyond close the system will purge cache immediately on the last close. With Buffer beyond close the system will retain buffers in storage longer. Because caching will be more efficient you will less I/O and less elapsed time but you may have more CPU because you'll be accessing data faster and because the cache overhead is larger with a larger cache. It's a good news story. Just be advised that you are going to have higher CPU utilization.

If you have a concern about how long data and directories are cached, you can change the LRU timing with the two parameters LRUTIME and LRUCYCLES

LRUTIME x LRUCYCLES = Amount of time that buffers stay in the cache. (The default for the LRU algorithm to process every 60 seconds and to keep a buffer for 15 processing cycles for a total of 15 minutes.)

Once buffer beyond close has been invoked either through PARMLIB or operator command, directory pages and optional member pages could be retained in cache for a longer duration.

PDSE LRU parameters are specifiable in the IGDSMSxx member of SYS1.PARMLIB.

PDSE LRU parameters are specifiable using the SETSMS command.

LRUTIME specifies the number of seconds (5 to 60) that the system waits between calls to the LRU routine. The LRU routine releases inactive buffers in 64-bit storage that are used to cache PDSE directory data and in the Hiperspace used to cache member data.

LRUCYCLES specifies the maximum number of times (5 to 240 cycles) that the LRU routine passes over inactive buffers before making them available for reuse.

Buffer beyond close: Validation

- Once buffer beyond close has been invoked either through PARMLIB or operator command, issue D SMS,OPTIONS command to ensure the line item is in effect.
- Refer to the SMF cache statistics:
  - SMF Type 14/15 with extended information segment type 6
  - SMF Type 42 subtype 1

Issuing D SMS,OPTIONS from the console will help you to know if the parameters for buffer beyond close are enable or disabled.

New types of SMF records have been added for PDSEs to validate the caching.

Refer to the z/OS® MVS System Management Facility (SMF), SA22-7630-12,

•Record Type 14 (0E). A new optional section called *PDSE Data Set Caching Statistics (type6)* will be added to the SMF Record Type 14 and 15 for PDSE data sets.

•Record Type 42 (2A) -- DFSMS™ Statistics and Configuration; subtype 1: BMF Totals Section

PDSE cache statistics in SMF 14/15

Extended information section descriptor word

Offsets		Name	Length	Format	Description
0	0	SMF14ESL	2	binary	Size of this extended information section (size of variable length fields including this 4 byte section descriptor word).
2	2		1	binary	Reserved
3	3	SMF14STY	1		Section type code. Type Meaning When Set 1 Compressed format data set section 2 SMS class information section 3 Step Information section 4 ISO/ANSI Version 4 CCSID (coded character set ID) information section 5 Additional data set characteristics section **6 PDSE data set caching statistics** 7 Key label data for tape encryption.

In addition, you now have new statistics that are provided in the SMF14/15 for the PDSE Subtype.

So there is a SMF 14/15 record and segment type 6 that is a PDSE segment. The contents of the PDSE segment type 6 provide caching statistics for the dataset.

PDSE cache statistics in SMF 14/15

PDSE data set caching statistics (type6)

This section describes the caching statistics for a PDSE data set.

Offsets		Name	Length	Format	Description
4	4	SMF14DRD	4	binary	Directory read request count
8	8	SMF14DRDH	4	binary	Directory read hit count
12	C	SMF14MRD	4	binary	Member read request count
16	10	SMF14MRDH	4	binary	Member read hit count
20	14	SMF14MCE	4	binary	Member Cache Eligible
24	18	SMF14MST	4	binary	Member Cache Stolen
28	1C	SMF14MNC	4	binary	Member Cache Eligible but not cached
32	20	SMF14MCF	4	binary	Member Cache Eligible but cache ull

You can get more information looking at the SMF 14/15 subtype 6 fields.

1. SMF14DRD: How many pages to try to read from directories

2. SMF14DRDH: How many of those pages were actually found in a Cache

3. Member counts has to do with Hiperspace caching. How many pages to try to read in the Hiperspace cache

4. How many were found in Hiperspace.

5. How many of them were eligible to be in Hiperspace

6. How many of them were found that had been stolen (indication that your Hiperspace is not big enough)

7. How many were eligible to be cached, but cache was full and they could not be put in cache (writing).

This information will help you to tune your system.

For directory cache is a combination of auxiliary and real storage.

For the Hiperspace is real storage. There is no auxiliary storage involved for the members.

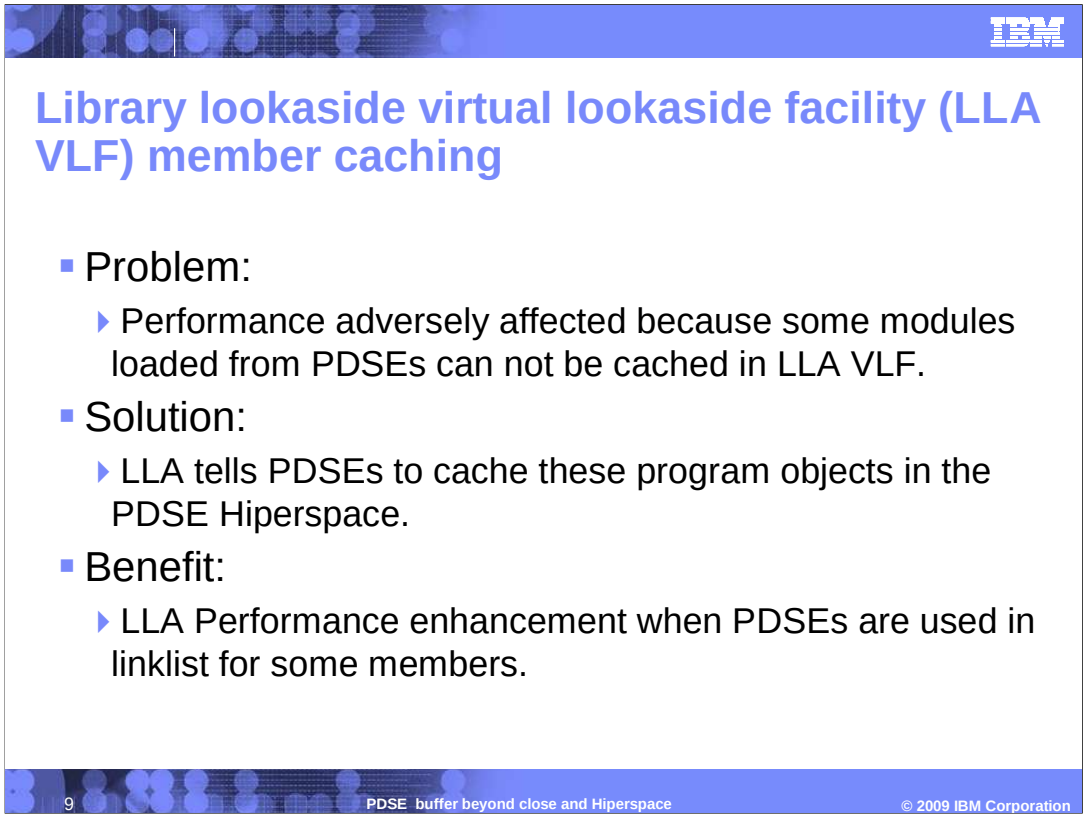

Read the Problem

The performance of some applications can be adversely affected because there are modules loaded from PDSEs which can not be cached in LLA VLF.

Read the solution

LLA tells PDSE support to cache these modules even though the datasets they are in are not marked as Hiperspace / member cache eligible

Read the Benefit.

This support will be active only if the IGWSMSxx parameter PDSE_HSP_SIZE(nnn) is non-zero. If Hiperspace size is zero, there is not going to be member caching.

References

- Partitioned Data Set Extended Usage Guide, Redbooks®
  - SG24-6106-01 or http://www.ibm.com/redbooks
- z/OS MVS System Management Facility (SMF)
  - SA22-7630-12
- z/OS DFSMS: Using Data Sets
  - SC26-7410-01

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_V1R0-PDSE-BufferBeyondClose.ppt

This module is also available in PDF format at: ../V1R0-PDSE-BufferBeyondClose.pdf

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DFSMS Hiperspace Redbooks z/OS

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. 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. 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.

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 NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

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.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. 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 ratios stated here.