

z/OS® large stand-alone dump Best practices

© 2007 IBM Corporation

This presentation summarizes the current best practices for taking and handling large stand-alone dumps in z/OS. This information is critical for speedy problem resolution in large memory environments. It is applicable to installations running at z/OS 1.6 and above.

Introduction

- This presentation discuss highlights of the z/OS large stand-alone dump best practices white paper (version 2):
 - ▶ Stand-alone dump planning
 - ▶ IPCS performance considerations
 - ▶ Post stand-alone dump processing
 - ▶ Getting IBM Support to analyze the dump
- Read the white paper after reviewing this presentation for further details and additional information:
<http://www.ibm.com/support/techdocs/atsmastr.nsf/Webindex/TD103286>

The highlights of the 'z/OS Best Practices: Large Stand-alone Dump Handling Version 2' white paper is discussed in this module. This paper, published in August 2007, describes a set of best practices for optimizing stand-alone dump data capture, post stand-alone dump processing and problem analysis.

Since this presentation does not cover all the details of the paper, it is strongly recommended that you also read the paper for further details and information, after viewing this presentation.

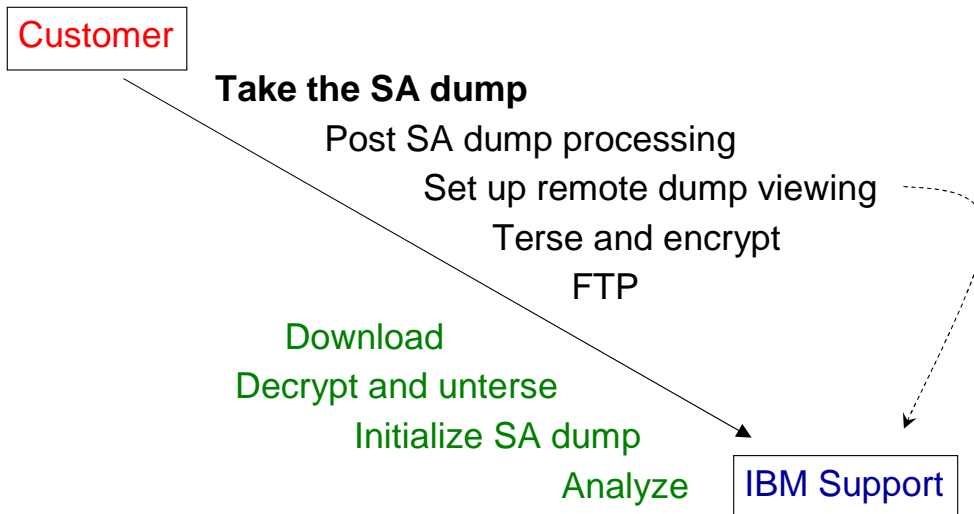
Agenda

- Stand-alone dump planning
- Stand-alone dump processing
- Post stand-alone dump processing
- Remote dump viewing
- Sending dump to IBM Support

This module will cover how to get the best stand-alone dump performance for large dumps, and the best ways of processing the dump so that it can be analyzed by IBM Service quickly and effectively. These topics will be discussed:

Planning for the stand-alone dump dataset, stand-alone dump processing, things to do after the stand-alone dump is taken, getting IBM Service to analyze the dump as soon as possible, and sending the dump to IBM Service.

Handling of a stand-alone dump



This chart demonstrates that there are a few steps that need to be taken before the stand-alone dump can be analyzed by IBM Service. For a very large stand-alone dump, some of these steps may take a long time. Speeding up any one of these steps will allow IBM Service to diagnose the problem sooner. Using a Remote Dump Viewing facility will enable IBM Service to review the dump while some of these steps are in progress.

After the stand-alone dump is taken, there is post-dump work such as copying the dump and preparing it for IPCS analysis. Then it would have to be tersed or encrypted, before FTP to IBM. IBM Service would need to download the dump, decrypt and unterse it, then initialize it under IPCS, before entering any IPCS commands to begin reviewing data from the dump.

SA dump planning

Strategy:

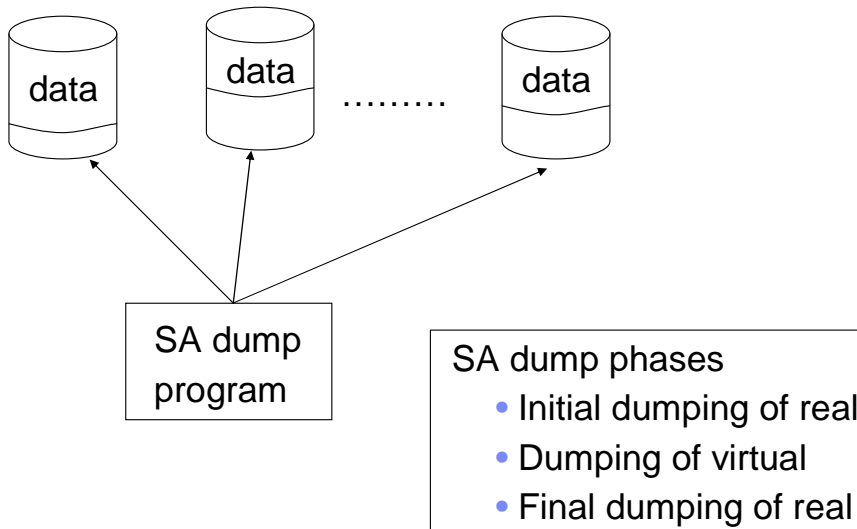
I/O technology is advancing proportionally with real storage size.
Maximize I/O rates by increasing parallelism and throughput.

- To obtain best dump performance
 - ▶ Use a multi-volume DASD SA dump dataset
 - ▶ Dedicate as many DASD volumes as possible (up to a maximum of 32 in z/OS 1.7 and above)
 - ▶ Use DASD volumes on separate LSS (or control units)
 - ▶ Use high performance DASD volumes
 - ▶ Configure DASDs to avoid other traffic affecting SA Dump I/O

Before z/OS 1.8, the operating system supports 128 gigabytes of real storage. This limit has been increased to 4 terabytes in z/OS 1.8. Stand-alone dumps will get bigger with more memory available in the system. Since I/O technology advances proportionally with size of real storage, the best stand-alone performance is realized when the dump is taken to a multi-volume dump data set on high performance DASD volumes. Configure each volume on a separate Logical Subsystem (LSS) or control unit to ensure maximum parallel operation. Dedicate a minimum of 4 to 5 DASD volumes for the dump data set. z/OS supports up to a 16-volume stand-alone dump data set through z/OS 1.6 and up to 32-volume in z/OS 1.7 and above. Configure, if possible, the DASD volumes to minimize the occurrence of other I/O activity such as database recovery affecting stand-alone dump I/O.

Be sure to catalog your stand-alone dump data set to prevent the 'wrong version' of the data set used for IPCS processing. Test your stand-alone dump procedures from time to time.

SA dump processing



A multi-volume stand-alone dump data set actually spread across all of the specified volumes, not each volume in succession. Data is written in parallel to all the volumes. The amount of data written on each volume varies because it is dependent on the I/O response of the volume. The order of how the data is written on the volumes is also dependent on the I/O response. After the dump completes, a merged version of the entire multi-volume dump is needed for IPCS analysis. To obtain the best IPCS performance, the merged version should contain records in logical order. This will be explained further on the next page.

There are three phases of stand-alone dump: initial dumping of real, dumping of virtual and final dumping of real. As you can see, dumping of virtual starts before all real has been dumped. This will allow the dump to contain more virtual storage in case it is stopped before all storage has been dumped.

Post SA dump processing

Strategy:

- Provide IBM Support access to the dump as soon as possible.
- Maximize IPCS performance upon processing the dump.

- Submit 2 IPCS COPYDUMP jobs in parallel
 - ▶ produce a merged copy of the SA Dump
 - ▶ produce a subset of the SA Dump (first 20 ASIDs)
- Ensure output of COPYDUMP is DFSMS-stripped
- Set up Remote Dump Viewing
- FTP the subset dump to IBM initially
- Use an IPCS Dump directory with the right attributes

After the stand-alone dump is taken, it is recommended that you use IPCS COPYDUMP, and not IEBGENER, to produce a merged version of the entire multi-volume dump. COPYDUMP usually takes less time than IEBGENER, and produce a merged dump with records in logical order. This will maximize IPCS performance when the dump is analyzed using IPCS. IEBGENER produces a merged dump with records in random order and will not give the same benefit. In order to allow IBM Service to start investigating the dump as soon as possible, a separate IPCS COPYDUMP job should be run in parallel to produce a subset dump of the first 20 ASIDs in the system. This subset dump is smaller and will take less time to be FTP to IBM initially. In many cases the subset dump can enable IBM Service to identify the failing product or component. Do review the white paper to get details on how to produce the subset dump.

IPCS performance is also improved if the output data set from IPCS COPYDUMP is DFSMS-stripped. IPCS uses a dump directory to maintain information about the layout and content of the dump. You can improve IPCS performance by specifying an optimum RECORDSIZE parameter to avoid CI (Control Interval) splits. Do review the white paper for more details.

Remote dump viewing

- RSVSF (Remote Screen Viewing Facility)
 - ▶ available as part of IBM Service Offerings
 - ▶ allows IBM support personnel to view a dump on your system through a session on your workstation

- AOS (Assist on Site)
 - ▶ available at <http://www.ibm.com/support/assistsite>
 - ▶ IBM Support can view your desktop and share control

Remote access products allow you to permit IBM Service to immediately log into an IPCS session on your system to analyze the dump with no initial data transfer. This should always be the first option in a time-critical situation.

In the United States there are two remote access products that can be used. RSVSF is available as part of IBM Service Offerings and allows IBM support personnel to view a dump through a session on your workstation. AOS is available at the web site mentioned and allows IBM Support personnel to access the dump by viewing your desktop and sharing control. Visit the AOS web site to read about how to use the tool so that you are familiar with it well before any critical situation arises.

Sending dump to IBM

- Terse
 - ▶ Use TRSMMAIN to compress data
 - ▶ TRSMMAIN is shipped as part of z/OS in 1.9 (AMATERSE)
- Encryption
 - ▶ For example, using secure FTP or IBM Encryption Facility
 - ▶ Place decryption information in the PMR
- FTP
 - ▶ PUTDOC allows dump to be split into multiple files and then sent by FTP
 - ▶ no improvement on FTP speed since files are sent serially, but will help recovery of data in case of network failure

To speed up the FTP process, the stand-alone dump should be compressed using TRSMMAIN before sending the data. In z/OS 1.9, TRSMMAIN is shipped as AMATERSE which is now part of z/OS.

You may also want to encrypt the data before tersing it and FTP. Remember to put the decryption information in the PMR.

When sending large dumps to IBM, using the PUTDOC facility allows the dump to be split into multiple files and then sent by FTP. By sending a large dump in smaller segments, recovery of data in the event of a network failure is limited to the segment in process.

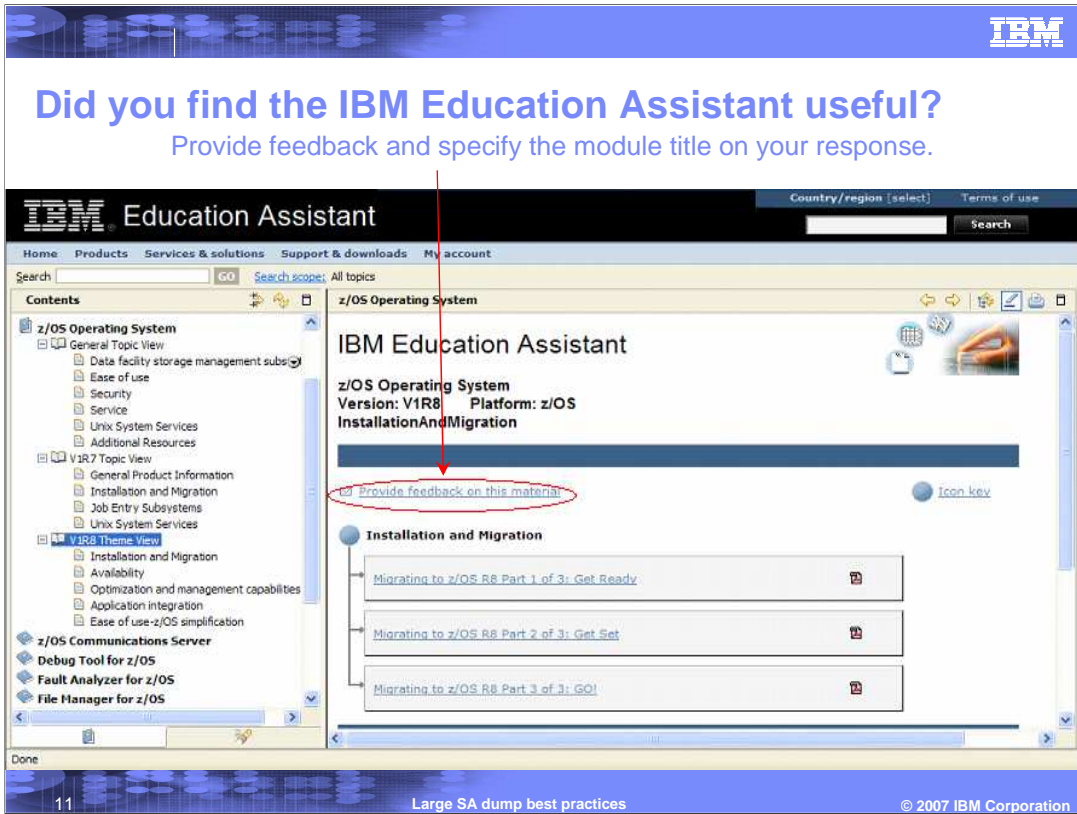
Review the white paper for web site information related to TRSMMAIN, IBM Encryption Facility and PUTDOC.

Summary

- Define the stand-alone dump as a multi-volume dump data set on high performance DASDs
- Configure the DASDs for maximum throughput
- Ensure the stand-alone dump data set is cataloged
- Submit 2 IPCS COPYDUMP jobs in parallel to obtain a subset dump and a merged dump
- Ensure output of COPYDUMP is DFSMS-stripped
- Use an IPCS dump directory with the right attributes
- FTP the subset dump to IBM initially
- Be familiar with how to set up remote dump viewing beforehand
- Terse dumps before FTP
- Use PUTDOC for large dumps
- Test stand-alone dump procedures

This presentation is intended to give you information about the best practices for large stand-alone dumps.

This chart gives you a summary of the important points that was covered. These steps should help in getting a quick resolution of your problem when a large stand-alone dump has to be taken.



In your feedback to IBM, answer these three questions:

How helpful was this presentation? Give it a rating from 1 to 5 where 1 = very helpful and 5 = not at all helpful

Did this presentation save you a service call to IBM? Yes or No

What other topics would you like to see covered?

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM z/OS

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.

Information is provided "AS IS" without warranty of any kind. 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.

© Copyright International Business Machines Corporation 2007. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.